IBM® Tivoli® IntelliWatch® User's Guide

Version 6.00.27.40

Document Number SC23-8877-00

November 2007

Table of Contents

List of Figures

List of Tables

Chapter

1 Introduction

Welcome to the IBM® Tivoli® IntelliWatch® User's Guide, your road map to improved Domino® server performance and reliability.

1.1.0 GETTING STARTED

1.1.1 Assumptions

This User's Guide is for Notes® Administrators and Application Developers, and makes the following assumptions:

- You have a thorough knowledge of the Lotus Notes® /Domino version supported on your platform.
- Lotus Domino is properly installed on your server.
- You have a thorough understanding of the applicable operating system(s).

1.1.2 Minimum System Requirements

Table 1-1 lists the minimum system requirements for the supported Windows® and UNIX® platforms. Please consult this table before installing IntelliWatch. Please Note...While many customers run IntelliWatch on non-English systems, it is only certified using English versions of Notes/Domino and the relevant operating system. Therefore, some features (e.g. PM NT Counters on German systems), may not function correctly out of the box.

- OS-level clustering is not supported.

Table 1-1. Minimum/Recommended System Requirements

System Requirements for version 6.00.27.36 and above (minimum supported/recommended) (For the most recent system requirements and version support information, go to www.ibm.com.)			
System Parameter	**Windows 2000 / 2003**	**AIX®**	**Solaris**
System Memory: running Notes Server	128MB / 256MB	256 MB per Domino partition	256 MB per Domino partition
System Memory: running Notes Client only	64MB / 128MB	N/A	N/A
Processor Speed	233MHz / 400MHz	Sufficient processor speed to run the Domino Server	Sufficient processor speed to run the Domino Server
Registry access	full access to the registry tree on NT: HKEY_LOCAL_MACHINE\SOFTWARE\Candle	N/A	N/A
Replication Check UI (stand-alone client only)	Windows 2000 (SP4)	N/A	N/A

Table 1-2. Additional Supported Platforms and Domino Versions

IntelliWatch Version	Platform Version	Lotus Domino Version
IntelliWatch 27.36 Server	Windows 2000 (SP 4); AIX 5.1 and 5.2; Solaris 8 and 9	Through 6.5.2
IntelliWatch 27.37 Server	Windows 2000 (SP 4); AIX 5.1 and 5.2; Solaris 8 and 9	Through 6.5.3
IntelliWatch 27.37a Server	Windows 2000 (SP4), Windows 2003 (SP4)	Through 6.5.4
IntelliWatch 27.37b Server 27.37b differs from 27.37a only in two binaries: niwagent.exe and iwnrt.dll. To see indications for installing 27.37b and to download a zip file containing these libraries, see http://www-1.ibm.com/support/docview.wss?uid=swg24011040.	Windows 2000 (SP4), Windows 2003 (SP4)	Through 6.5.4
IntelliWatch 27.37a Client	Windows XP, Windows 2000 (SP4), Windows 2003 (SP4)	Through 6.5.4
IntelliWatch 27.38 Server IntelliWatch 27.39 Server IntelliWatch 27.40 Server	**Linux® (Intel)** ■ **Through Domino 7.0.2 (where supported on OS by Lotus)** **Windows 2000 (SP4) and Windows 2003 (SP4); AIX 5.1, 5.2, 5.3; Solaris 8, 9, and 10** ■ **Through Domino 7.0.2 (where supported on OS by Lotus)**	
Note: Intelliwatch is supported on Domino 6.5.x, and was certified on Domino 7.x.y. and Domino 8.x. For additional information about the operating systems that are supported, see http://www-306.ibm.com/software/sysmgmt/products/support/Tivoli_Supported_Platforms.html		

1.2.0 QUICK TOUR OF INTELLIWATCH

1.2.1 What's new?

Many of the following new features have been in IntelliWatch since the GA release of version 6.0. Other features have been added in a later maintenance release. Since a few customers are still running Pinnacle 99, and may be evaluating when move to the Enterprise version, all new Enterprise-level features are listed below.

For the benefit of those customers already running the Enterprise version, features added in later maintenance releases include the number of that release in parentheses.

- **a more responsive Console**
  - via a browser

    Console.nsf loads in Internet Explorer, offering convenient access—without requiring you to install software on the workstation.
  - via an IntelliWatch stand-alone client

    The stand-alone client installs on your Admin workstation. In addition to the full range of functionality available via the browser client, the stand-alone client offers:
    - Replication Check utility (see *11.4.0.0*)
    - Option of selecting the initial IntelliWatch Solution that will be displayed when the Console loads.
- **security mechanism**
  - via a browser

    Browser-level check

    ACL check of console.nsf and iwasw.nsf
  - via the stand-alone client

    ACL check of console.nsf and iwasw.nsf
- **IW Message Center Gateway**

More flexible messaging support has been added to IntelliWatch by means of the IW Message Center Gateway.

The Gateway runs on both NT and AIX, and receives messages over TCP/IP. Based on internal variables, the Gateway determines both message type and content, then relays the notification in a form appropriate to the recipient/application, including:

- Tivoli Event Console
- SNMP Network Manager
- NT Event Log
- Paging service (NT only)

IwMCG requires that TCP/IP be available on both the client and server systems.

- **new Monitor Command types**
  - Tivoli Tec Event

    IntelliWatch Triggers can now send alerts to a Tivoli console.
  - ACL Changer

    Automate changes to the Access Control List of Notes databases using IntelliWatch Triggers and this new Command type.
- **improvements to Analyzer**

The Analyzer server task has been made more robust, resulting in a significant increase in the speed of statistic processing and report creation.

- **Monitor Triggers on Blackberry servers**

Though no product changes were involved, we're please to pass on the news that some customers have been using IntelliWatch Monitor Triggers to monitor Blackberry servers—starting with Pinnacle 99!

We hope to add some default Triggers for Blackberry, QuickPlace and Sametime in future maintenance releases of IntelliWatch.

- **new IntelliWatch <KEYWORD> (27.33)**

<DATABASE_OCCURRENCES_FOUND> returns the number of times the search string was actually found (as opposed to the number stated on the Condition tab of the Trigger).

- **ASW: Escalated Action Profiles (27.33)**

New in maintenance release 27.33 are escalated Action Profiles, that allow admins to take increasing aggressive measures, the longer a server is unresponsive.

- **ASW: "I am alive" messages from Monitor (27.36)**

Starting in maintenance release 27.36, each time a Trigger evaluates, a special "I am alive" message creates (or updates) a document in **iwasw.nsf** which contains the server name and the date.

Thanks to this feature, you have the option of being notified in situations where the Domino server is up, but where the IntelliWatch IWAGENT task may have hung.

- **Monitor: New optional repository for Availability Report documents (27.36)**

Starting with maintenance release 27.36, the option exists of using an alternative database to store the Server Status Log on which Monitor's Availability Reports are based, namely **iwstatus.nsf**.

*The name **iwstatus.nsf** is hardcoded; changing it is not supported.*

1.2.1.1 Why this new database?

Occasionally, customers have experienced errors in Monitor's Availability Statistics, due to a database purge interval that was shorter than their maintenance interval.

Consequently, Up and Down documents (in the Server Status Log View) got out of sync, resulting in errors.

- **Monitor: new <KEYWORD> (27.36)**

New in maintenance release 27.36 is the <ESCALATION_INDEX> keyword, which indicates the number of consecutive monitoring intervals during which a Trigger Condition has remained true.

1.2.2 Configuration

Most product configuration is done via the Pinnacle Console (running on NT at your Admin workstation).

The Pinnacle Console is available in two forms:

- Pinnacle Console via a browser

  A Notes database (**console.nsf**) residing on your Primary Server.

  Currently supported on Internet Explorer (version 4.0 or higher).

- Pinnacle Console as stand-alone client

Local client that installs on your Admin workstation.

You must install this version of the Console to use the Replication Check utility.

In addition, several features of IntelliWatch can be configured by accessing the IntelliWatch databases directly, using either a Notes client or Internet Explorer.

Certain features of IntelliWatch must be configured in the NT registry (or in IntelliWatch .ini files on UNIX® systems). Still other settings must be added to/edited in the Notes.ini. These configuration options are mentioned where applicable.

1.2.3 Chapters

1: Introduction

Overview of User's Guide, and summary of IntelliWatch main features.

2: Installation Guide

Step-by-step instructions for installing on NT.

3: Management Agents

Explains the comprehensive problem detection/correction capabilities offered by IntelliWatch Triggers and Commands.

Practical usage suggestions are provided for each Trigger and Command type.

4: Advanced ServerWatch

Advanced ServerWatch (ASW) shows you—at a glance—the connectivity status of all monitored Domino servers.

Monitor Triggers can send IWAlerts to the ASW Console, where you can view these up-to-the-minute warning and status messages.

The response time of server monitored by ASW can be logged to **iwstats.nsf**, for subsequent retrieval by Analyzer.

5: Performance Manager

Pinnacle Performance Manager (PM) allows you to manipulate Notes statistics, and to create custom statistics.

PM statistics can be imported into IntelliWatch Analyzer, to produce reports on all facets of Domino server activity.

Used as statistic thresholds in Monitor Triggers, IntelliWatch statistics allow you to fine-tune your environment—and to react to potential problems before they seriously impact Domino performance.

6: Tracer

Tracer helps diagnose bottlenecks caused by slow response times in Notes databases.

Events are recorded in real time, and stored in a log for processing. When displayed at the Pinnacle Console, Traces can be fine-tuned using filters.

Traces can also be saved for later analysis.

7: IntelliWatch Message Center Gateway

The sole function of the IntelliWatch Messaging Center Gateway is to relay messages received from IntelliWatch components

The application's simplicity makes for amazing flexibility. IwMCG receives a message over TCP/IP, and, based on internal variables, determines both message type and content. The Messaging Center then relays the notification in a form appropriate to the recipient/application.

8: Crash Detection

IntelliWatch Crash Detection monitors server availability locally, complementing the remote connectivity monitoring provided by Advanced ServerWatch.

The local monitoring functionality is based on the ability of Crash Detection to open a local database (the default is Log.nsf). If the database can be opened, the server is considered to be running; if the database cannot be opened, Crash Detection follows a series of steps (the exact sequence of which depends on various configuration options).

9: Analyzer

The report engine of IntelliWatch, Analyzer can import native Notes statistics, custom IntelliWatch stats, and response-time statistics logged by Advanced ServerWatch.

Analyzer reports can be scheduled, or run on demand. They are stored in a local Notes database, and can be distributed via Notes mail.

10: Configuring IntelliWatch via Notes

Explains how to configure the following IntelliWatch components via a Notes client:

- Management Agents
- Advanced ServerWatch
- Performance Manager
- Analyzer

11: Console Utilities

Several utilities are available from the IntelliWatch user interface:

- Parameter Configuration Utility
- Send Page Utility
- Remote Recycle Utility
- Replication Check (stand-alone client only)
- Send SNMP Trap Utility

12: Command-line Utilities

Effective remedies for common problems, IntelliWatch command-line utilities simplify many management tasks, from sending mail to recycling your Domino server.

13: Feature Enhancements

Overview of feature enhancements introduced in the latest maintenance release.

1.2.4 Appendices

Appendices provide supplementary information in the following areas:

A: Definitions of PM Statistics Variables

Field-by-field explanation of all PM Statistics variables.

B: IntelliWatch Keywords

Definitions for all IntelliWatch Keywords.

C: Paging Error Messages

List of Paging Error Messages by error code.

D: IntelliWatch Databases: Replication and Template Usage

Table of information detailing which IntelliWatch databases replicate (and which do not), as well as which templates are used in the creation of databases.

E: NT Setup Dialogs

Screenshots of all Setup dialogs displayed when IntelliWatch is installed on NT/Win2000.

F: IntelliWatch Configuration Wizard

Screenshots of all main Configuration Wizard dialogs.

G: Integrating Monitor 5.0 for AS/400 with IntelliWatch

Provisos for integrating Monitor 5.0 for AS/400 in a mixed environment with IntelliWatch.

H: Data Returned by PM Statistic Types

What is returned by each Performance Manager statistic type.

I: Ports Used by IntelliWatch

Useful information on the ports used by IntelliWatch, on both partitioned and non-partitioned Domino servers.

J: Getting More Help

How to obtain additional help in using IntelliWatch.

1.3.0 ACCESSING THE PINNACLE CONSOLE

New to IntelliWatch is a Web-based Console, as well as a more responsive stand-alone version.

Both Console versions have been enhanced with a security mechanism.

Basic access to each of the Console versions is obtained as follows:

PINNACLE CONSOLE IN A BROWSER

TO ACCESS THE PRIMARY SERVER:

1. Make sure that 1) the Primary Server is running, and 2) the HTTP task is loaded, and 3) is accessible via TCP/IP.
2. Launch Internet Explorer at your Admin workstation.
3. Type the URL of your Primary Server in the IE Address field, based on the following model:

```
http://[DomainNameOfPrimaryServer]/[IntelliWatchDataDir]/console.nsf
```

Depending on local configuration (and from where the Primary Server is being accessed), the complete

Internet domain name may be required for the DomainNameOfPrimaryServer.

For example: If the Host Name of your Primary Server is MyServer/MyDomain. From within the domain, "MyServer" should suffice as the DomainNameOfPrimaryServer. *Under some conditions, however, you may need to enter "MyServer.MyDomain.com".*

At this point, you are confronted with the browser level security check (for details, see *1.4.1*, below). Assuming this step is successfully navigated, the opening page of the database will be displayed in your browser, allowing you to proceed as follows:

4 Click the Pinnacle Console link in the upper left of the page.

At this point, you are confronted with the ACL level security check (see *1.4.2*, below). Assuming this step is successfully navigated, the Console will be displayed in your browser.

The first time you connect to the Pinnacle Console, you are informed that no Advanced ServerWatch Hubs exist, and asked if you would like to create Hubs at this time. Decline if you want to create Hubs at a later time, but bear in mind that, until at least one Hub is created, this dialog will be displayed each time the Console loads.

STAND-ALONE PINNACLE CONSOLE

The following procedure assumes you have already successfully run the Setup to install the stand-alone Console. If this is not the case, and you need assistance to install this version of the Pinnacle Console, please see *"IntelliWatch Stand-alone client" on page 37*.

TO ACCESS THE PRIMARY SERVER:

1 Make sure that 1) the Primary Server is running, and 2) the HTTP task is loaded, and 3) is accessible via TCP/IP.

2 Launch the stand-alone Console by going to **Start > Programs > IntelliWatch > Pinnacle Console**. The dialog in *Figure 1-1* will be displayed.

FIGURE 1-1: Initial-launch dialog

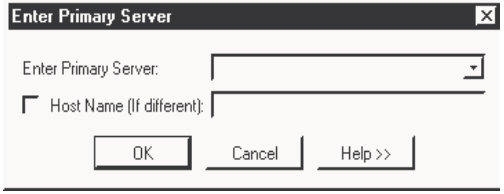

3 Enter the name of your Primary Server, or select it using the drop-down.

If the Host Name of the target system differs from the Domino Server Name, select the checkbox and enter the Host Name in the textbox provided.

4 Click OK.

At this point, you are confronted with the ACL level security check (see *1.4.2*, below). Assuming this step is successfully navigated, the Console will be displayed in your browser.

As with the web-based Console, if no Advanced ServerWatch Hubs have as yet been created, a dialog is displayed giving you that option. To proceed without creating Hubs, click No.

1.4.0 SECURITY MECHANISM

New to IntelliWatch is a security check that is performed whenever a User tries to connect to the Pinnacle Console.

The number of stages depends on the means used to connect to the server:

- via a browser
  - Browser-level check (*Figure 1-2*), followed by a check of the relevant ACLs (*Figure 1-3*).
- via the Pinnacle stand-alone client
  - ACL check only (*Figure 1-3*).

1.4.1 Browser level

- User Name and Password

The browser-level check requires you to enter a User Name found in the Notes NAB *on the target server*, as well as the *Internet* password associated with that User.

This Internet password is found/entered on the *Basics* tab of the relevant Person document.

FIGURE 1-2: Initial security check

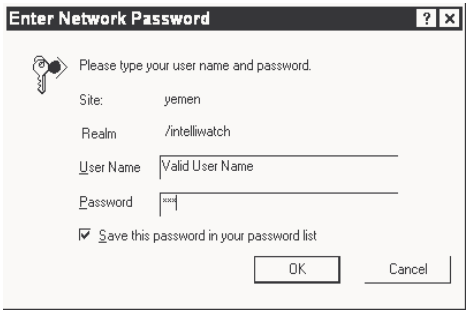

An optional checkbox allows you to save this password for future sessions (although this largely defeats its purpose).

If no Internet password is associated with the User attempting the connection, either enter one in the relevant Person document and reattempt the connection, or log in using a different User Name (for which there is such a password).

If either log-in value is incorrect, the connection is refused, and you are given another opportunity to enter correct information. (For a flow chart detailing the mechanism, see *Figure 1-4*.)

Assuming successful completion of this first log-in stage, the security mechanism proceeds with a check of the relevant ACLs.

1.4.2 ACL check

A User Name and Password must be entered (see *Figure 1-3*). The User Name will be checked against the relevant ACL(s); the Password, again, is the Internet password referred to under *1.4.1*, above.

FIGURE 1-3: ACL security check

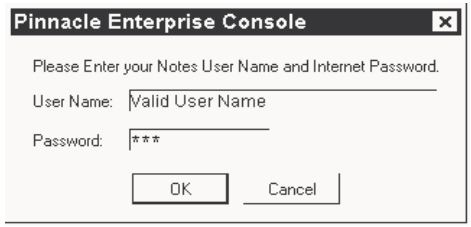

This need not be the same User Name/Password entered at the browser level.

- Which ACLs are checked?

The answer to this question depends on whether Advanced ServerWatch Hubs have already been created via the Pinnacle Console.

As they are created, Hubs are appended to the list of *Hub_Servers* in the **pinnacle.ini** file on the client machine. Each time the Console database is loaded from that server, the list of ASW Hubs to be displayed is obtained from this file.

- No Hubs yet created via Console
  - A search for the entered User Name is carried out in the ACL of **console.nsf**.
    - The entered User must have an Access Level of READER or above, or the log-in fails. (See *1.4.3*, below.)
    - The User may be part of a group.
  - Assuming the entered User Name and Password are correct, the Pinnacle Console now loads.
  - You are informed that no Hubs exist, and are asked if you want to create one.
    - If you say No.

      The ACL of **iwasw.nsf** is not checked.
    - If you say Yes.

      The Console attempts to connect to the Advanced ServerWatch task on the server you select, and the ACL check of that database is carried out.

      The User Name and Password entered in *Figure 1-3* are used to check the ACL of **iwasw.nsf**. If this User has the necessary access privileges, the connection succeeds. Otherwise, an error occurs, and you are prompted to enter another User Name and Password.

- ASW Hubs listed in **pinnacle.ini**
  - First, a search for the entered User Name is done in the ACL of the Pinnacle Console database (**console.nsf**).
    - The entered User must have an Access Level of READER or above, or the log-in fails. (See *1.4.3*, below.)
    - The User may be part of a group.
  - Assuming the ACL-check of the Console database is successful, the ACL-check of the Advanced ServerWatch database (**iwasw.nsf**) is carried out for each Hub listed.

– The User Name and Password entered in *Figure 1-3* are used to check the ACL of iwasw.nsf. If this User has the necessary access privileges, the connection succeeds. Otherwise, an error occurs, and you are prompted to enter another User Name and Password.

– If more than one Hub is listed, the same process repeats for each Hub.

In the event an alternate User Name must be used to gain access to a given Hub's database, the process will revert to the initial User Name entered for each subsequent Hub (the only one that is cached).

- What if either ACL check fails?

After a brief delay, you are given another opportunity to correct/change the User Name and Password entered. Bear in mind, however, that after the first few tries, the time interval between tries becomes progressively longer, to discourage nefarious (unauthorized) users.

- Where is the User Name stored?

After the initial log-in on a system, the User Name value is stored as *Last_User=* in the **pinnacle.ini** file on the client machine.

Note that the Password is not stored in this file, and must be entered manually each time the dialog is presented.

1.4.3 Access Levels

Access levels play a role in the functionality available from this point on.

IntelliWatch, uses the following four levels:

- NO ACCESS
  - Connection will be refused.
  - Enter a User Name and Password with the correct access level, or cancel the operation until one can be obtained.
- READER
  - Can access all Pinnacle Solutions, but cannot create, edit or delete documents.
  - Should you attempt to write to a database (saving a Profile in Advanced ServerWatch, for example), you are prompted for a User Name and Password with the appropriate access level.
- READ/WRITE
  - Full access—with the exception of the Remote Recycle Utility.
- MANAGER
  - Full access.

1.4.4 More Secure Password format

On Notes 4.6 and above, a "More Secure" format can be applied to Internet passwords.

This option is selected via the Actions menu. The item in question is:

Upgrade to More Secure Internet Password Format

There is no UNDO for this more secure encryption.

Users with passwords so encrypted need to add a Person (User) to the **console.nsf** database to be able to access the Pinnacle Console

The solution below only applies if the Users who will be logging in to the Pinnacle Console have Internet passwords encrypted using the "More Secure" format.

To determine this for a given User, do the following:

1. Open the relevant Person document in your Domino Directory.
2. Put the document into Edit mode.
3. Examine the encrypted Internet password (located on the Basics tab).

- if password is "More Secure"
  - It will consist of numbers and both UPPER- and lower-case letters, and perhaps other symbols.
- if the password is NOT "More Secure"
  - It will consist of numbers and UPPER-case letters only.

If the Internet passwords of all Users who will be logging in to the Pinnacle Console fall into the second category, you can ignore the following solution.

SOLUTION FOR USERS WITH "MORE SECURE" INTERNET PASSWORDS

1. Connect to the Pinnacle Console database on the target server via Notes or a browser.

This brings up the initial 'links' page of the Pinnacle Console.

2. Click on the Edit user password link on the opening page of the Console database.
3. Click on the Add Person icon, and complete and save the displayed form.

This log-in ID is stored in the console.nsf database, and will function as a sort of 'mini-NAB'.

Now, when a User logs in, User Names/passwords are checked in the following order:

- Console.nsf checks the NAB for the User Name entered.
- If the User Name IS found in the NAB, the password is checked for "More Secure" encryption.
  - If the password does NOT use "More Secure" encryption, it is checked for validity. If the password is invalid, the log-in fails, and the User is prompted to try again.
  - If the password DOES use "More Secure" encryption, console.nsf checks its internal 'mini-NAB' for the User Name/password entered.
    - If the User Name/password DOES exist in console.nsf's internal 'mini-NAB', the log-in succeeds.
    - If the User Name/password does NOT exist in console.nsf's internal 'mini-NAB', the log-in fails, and the User is prompted to try again.
- If the User Name is NOT found in the NAB, the log-in fails.

If your Internet password does not use the More Secure format, these documents are not used by the Console.

FIGURE 1-4: Security mechanism flow chart

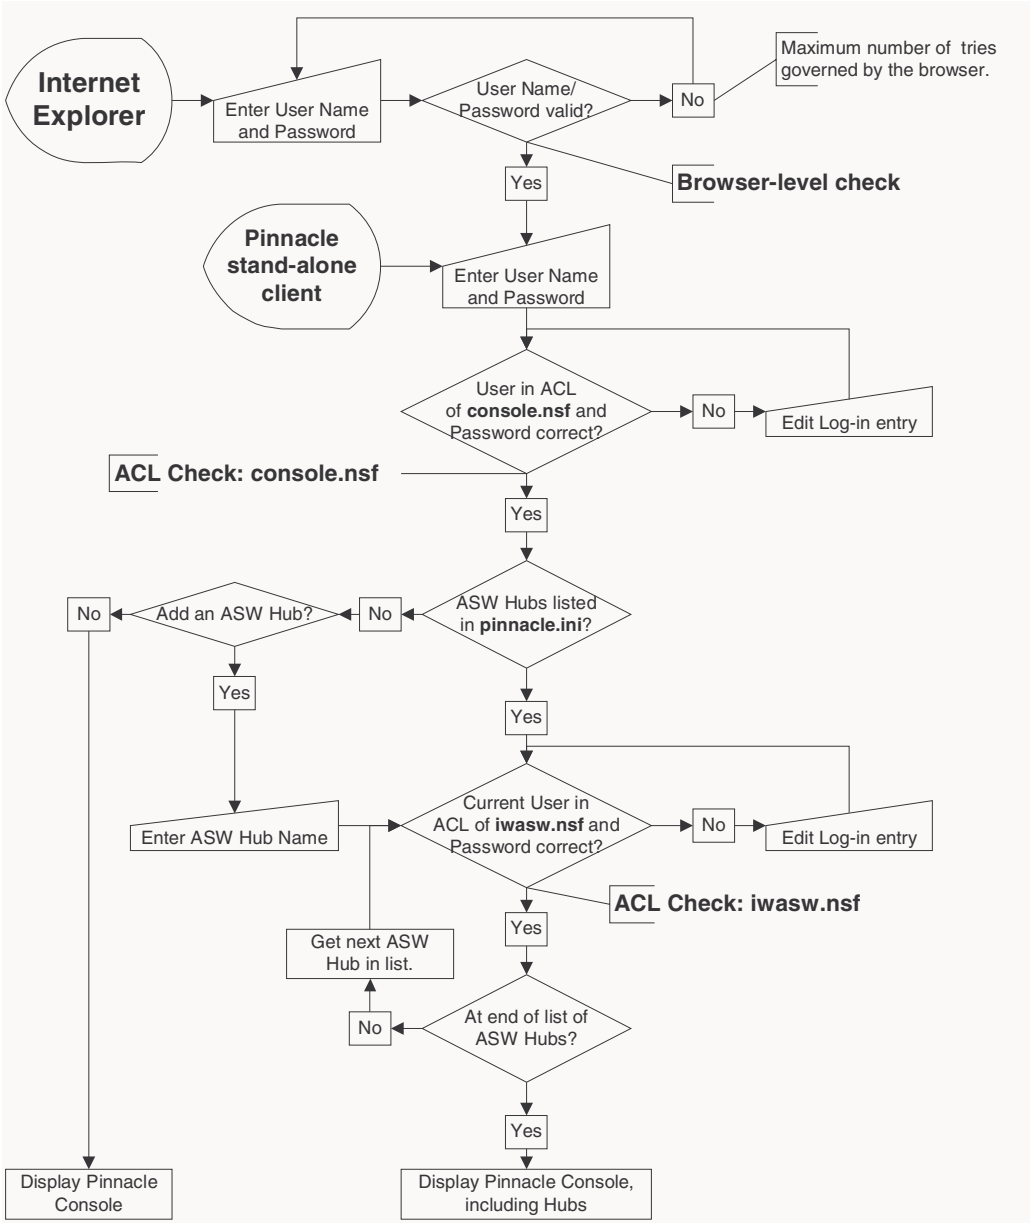

1.5.0 CROSS-REFERENCES

Cross references in the User's Guide that incorporate numbers/letters are to be interpreted as follows:

- first character (from left)
  - Chapter number/Appendix letter
- second digit (from left)
  - Level 1 Heading
- third digit (from left)
  - Level 2 Heading
- fourth digit (from left)
  - Level 3 Heading

1.6.0 NOTE TYPES

The detective signals text that provides additional details as to product functionality.

A light bulb signals a Hint as to how you can most effectively use the product.

Traffic light signals text of a cautionary nature, often concerning actions that, while possible, are not recommended.

Remember ...

Text to remind you of a detail covered previously.

1.7.0 ADOBE PORTABLE DOCUMENT FORMAT

IBM supplies documentation in the Adobe Portable Document Format (PDF). The Adobe Acrobat Reader prints PDF documents with the fonts, formatting, and graphics in the original document. To print a IBM document, do the following.

1. Specify the print options for your system. From the Acrobat Reader Menu bar, go to **File > Print Setup...** and make your selections. A setting of 300 dpi is highly recommended as is duplex printing if your printer supports this option.
2. To start printing, go to **File > Print** via the Acrobat Reader drop-down menus.
3. On the Print pop-up, select one of the **Print Range** options for
   - a single page
   - a range of pages
   - all the document
4. (Optional). Select the Shrink to Fit option if you need to fit oversize pages to the paper size currently loaded on your printer.

1.7.1 Printing problems?

The print quality of your output is ultimately determined by your printer. Sometimes printing problems can occur. If you experience printing problems, potential areas to check are:

- settings for your printer and printer driver. (The dpi settings for both your driver and printer should be the same. A setting of 300 dpi is recommended.)
- the printer driver you are using. (You may need a different printer driver or the Universal Printer driver from Adobe. This free printer driver is available at www.adobe.com.)
- the halftone/graphics color adjustment for printing color on black and white printers (check the printer properties under **Start > Settings > Printer**). For more information, see the online help for the Acrobat Reader.
- the amount of available memory in your printer. (Insufficient memory can cause a document or graphics to fail to print.)

For additional information on printing problems, refer to the documentation for your printer or contact your printer manufacturer.

1.7.2 Ordering additional product documentation

To order additional product manuals, contact your IBM Software Support representative.

Installation Guide

Chapter

2

2.1.0.0 Configuring IntelliWatch for your environment

Notes topologies differ—sometimes quite markedly—from one company to another. Still, certain guidelines apply to most environments where IntelliWatch is used to monitor Domino servers. We offer the following sample architecture to help you get started.

Chapter Contents

2.1.1.0 Sample architecture

The XYZ Company, Inc. has a 40-server Notes environment, divided into two server groups of 20 systems each.

An effective IntelliWatch configuration would include:

- one Primary Server

  The Master Set of those IntelliWatch databases that replicate resides here. Edit these databases here, then replicate the changes to your other servers.

  IntelliWatch components to install:

  – IntelliWatch Monitor
  – Pinnacle Performance Manager
  – Tracer Server

- two Advanced ServerWatch Hubs

  In this example, each Hub watches a group of twenty servers, 19 Managed Servers plus the ASW Sub-Hub.

  IntelliWatch components to install:

  – Advanced ServerWatch Server
  – IntelliWatch Monitor
  – Pinnacle Performance Manager
  – Tracer Server

- one Analyzer Server

  Statistics (both Notes and IntelliWatch) must be replicated to this system from all servers you want to include in reports.

  IntelliWatch components to install:

  – Analyzer Server
  – IntelliWatch Monitor
  – Pinnacle Performance Manager
  – Tracer Server

- Balance of systems as Managed servers

  This category includes your Application and Mail servers.

  IntelliWatch components to install:

  – IntelliWatch Monitor
  – Pinnacle Performance Manager
  – Tracer Server

- Admin workstations (NT)

  IntelliWatch components to install:

  – Messaging Center Gateway (on one system)
  – IntelliWatch stand-alone client, if desired
  – Tracer Client

2.1.2.0 Installation guidelines

This example illustrates the following principles of sound IntelliWatch architecture:

- Centralized database management

  When installing IntelliWatch on other systems, replicate databases from the Primary Server (bearing in mind that not all IntelliWatch databases replicate—iwasw.nsf should not, for example.

  Making changes only in the Master Copy on the Primary Server virtually eliminates the chance of replication conflicts.

For product trials on test systems, we suggest the "Use Default Triggers" setting (for the selection dialog, see Figure E-19 on page 418). On production systems, we recommend the "Disable All Triggers" option.

- Monitor Advanced ServerWatch Hubs

  The smallest Notes environment should run a minimum of two ASW Hubs.

  Why is this necessary, given the fact that a single ASW Hub can effectively monitor as many as 50 servers?

  Because with only a single instance of ASW Server in your environment, should the Hub system experience difficulties, monitoring of all spoke servers stops until you correct the problem on the Hub.

  (See, for example, the diagrammed configuration at *Figure 4-1, "ASW in a seven-server environment," on page 157*.)

Although each partition of a Domino partitioned server can theoretically run its own instance of the ASW server task, this is not recommended.

Should a partitioned machine running multiple instances of the ASW task fail, usually all instances will cease to monitor their assigned servers.

We recommend installing only a single instance of Advanced ServerWatch on a given system.

- Centralize reporting using Analyzer

By collecting statistics centrally, reports can be as comprehensive as you like, since data for all servers is available on a single system.

Analyzer retrieves statistics from the local copy of Notes and IntelliWatch statistics databases.

2.1.2.1 Should I use the Pinnacle Configuration Wizard?

We do not recommend running the Pinnacle Configuration Wizard when installing IntelliWatch.

The Wizard was designed to be used by on-site SEs, to automate product demonstrations—not to be used for tailoring IntelliWatch to the needs of a customer's production environment.

Running the Wizard as part of a normal installation creates a significant likelihood that you will want (or need) to disable (or modify) much of what the Wizard generates. (See also *"IntelliWatch Pinnacle Configuration Wizard" on page 441*

2.1.2.2 On-site Services

If you need more detailed assistance in setting up IntelliWatch in your environment than you find in this user's guide, we invite you to contact a sales representative, who can explain the availability and cost of on-site services.

2.1.3.0 Installation procedures

To install IntelliWatch, run the self-extracting Setup file, and respond to the dialogs displayed (command-line questions on UNIX systems). Which dialogs/questions

are displayed depends on the choices you make as the Setup proceeds.

For the location of the setup files on NT and UNIX systems, as well as platform-related differences in procedure, see the individual sections, below.

2.1.3.1 Getting help during the installation

Which setup dialogs/questions are displayed depends on the choices you make as the installation proceeds. For assistance with a particular setup dialog, see the Appendix *"NT Setup Dialogs" on page 399*. While the presentation of installation steps on UNIX systems differs from the NT dialog method, the information to be entered is, for all practical purposes, identical. Therefore, no separate section is provided for the steps of the installation process on UNIX systems.

Dialogs/questions that appear only in connection with the Upgrade are discussed after the main series. A series of links is provided to lead you through the dialogs displayed during an upgrade on NT (see *2.1.3.8*, below).

2.1.3.2 NT 4.0/Windows 2000

Start the setup by launching the installation file located in the **windows** folder on the IntelliWatch CD. The file is called **wpv06***.exe--the asterisks represent characters that vary by maintenance level. Alternatively, open index.html (on the root of the CD) in Internet Explorer, and click on the Copy Installation File link. You are then given the choice of copying the file to the local system (the default option), or of running the file from its current location.

Accepting the default requires that you subsequently navigate to the file's location on your local system to launch the setup. Run the file from the CD to install immediately.

The first time you run **wpv06***.exe**, installation files are unpacked to a default directory on the local system (the location varies by OS version). Among the unpacked files is **setup.exe**, which launches automatically once files are unpacked.

*If the installation process is aborted after unpacking has completed, the setup can be run at a later time by launching **setup.exe** directly. You need not go through the unpacking process a second time.*

2.1.3.3 NT Setup dialogs

Which setup dialogs are displayed depends on the choices you make as the installation proceeds. For assistance with a particular setup dialog, see the Appendix *"NT Setup Dialogs" on page 399*.

Dialogs that appear only in connection with the Upgrade are discussed after the main series. A series of links is provided to lead you through the dialogs displayed during an upgrade.

2.1.3.4 UNIX systems

The setup file for your UNIX system is located in one of two subfolders of the **unix** folder on the CD: either aix** or sol**. Copy the file appropriate to your operating system to a location of your choice.

Alternatively, on a system supporting a browser, open index.html (on the root of the CD) in Internet Explorer, and click on the Copy Installation File link for your UNIX platform. Select the default option of copying the file to the local system.

Before running the setup, you must first unzip, then untar, the installation file.

From the directory where the file was untarred, issue the following command on both Solaris and AIX

```
./setup
```

*To install IntelliWatch successfully, you must be logged on as **root**.*

2.1.3.5 On Solaris systems running Domino 6 and above

Consult the readme.txt file, found in the folder where you tarred the installation files.Un-installing IntelliWatch

Should you want to un-install IntelliWatch IntelliWatch, follow these simple steps.

2.1.3.6 On Windows

1. Go to **Start > Settings > Control Panel**.
2. Double-click the **Add/Remove Programs** icon, bringing up a dialog that allows you to select programs.
3. Highlight IntelliWatch, and follow the instructions on the ensuing series of dialogs.

2.1.3.7 On AIX systems

1. Log in as **root**.
2. Issue the **./uninst** command from the folder where the untarred IntelliWatch setup files are located.

*If that folder no longer exists, or installation files have been removed, untar the setup to a directory of your choosing, and run **./uninst** from that location.*

3. Select the IntelliWatch installation to be removed from the list displayed.

If this is a non-partitioned system, or if you are removing the last of a series of IntelliWatch partitioned installations, you will be asked if you want to remove program executables as well.

2.1.3.8 Supported upgrade path

When upgrading from Pinnacle 99, only one upgrade path is supported:

- from the latest version of Pinnacle 99 to IntelliWatch
  - on NT, this is version 21.10
  - on UNIX systems, this is version 22.14

Upgrading from earlier versions of Pinnacle 99, or from IntelliWatch versions prior to Pinnacle 99 is not supported.

To upgrade Pinnacle 99 to IntelliWatch, go to *page 399*.

2.1.3.9 Should I run Updatexxx.exe, or the Full Install?

With the advent of IntelliWatch Pinnacle for the Enterprise, the latest maintenance release is no longer included in the Full install. What does this mean, in practical terms?

- to update IntelliWatch v. 6.0 on Windows

  To bring a pre-existing installation of IntelliWatch Pinnacle for the Enterprise up to the latest maintenance release on Windows systems, run the Update[xxx].exe, NOT the full install.

- to upgrade from Pinnacle 99 (v. 5.0) to the Enterprise version on Windows

  To upgrade Pinnacle 99 to the Enterprise version of IntelliWatch, run the full install.

- on UNIX systems

  Whether updating a pre-existing installation of IntelliWatch Pinnacle for the Enterprise, or upgrading a system from Pinnacle 99 to the latest maintenance release of IntelliWatch, run the full install.

2.1.4.0 When you upgrade Domino

From time to time, customers will run incremental upgrades for their Domino installations, rather than doing fresh installs. How does this affect IntelliWatch installations on those systems?

IntelliWatch runs within the Domino envelope, so to speak, and updating your Domino version without un- and re-installing IntelliWatch can lead to unpredictable behavior of IntelliWatch components or features.

Although no instances of IntelliWatch-induced server crashes resulting from upgrading Domino in this way have been reported, that possibility cannot be ruled out.

Therefore, updating your Domino version without un- and re-installing IntelliWatch is not supported.

Preserving your IntelliWatch settings under these circumstances is not difficult, however. See the following section.

2.1.4.1 Preserving settings when Un- and Re-installing IntelliWatch

On rare occasions, such as when running an incremental install on your Domino installations, you may find it necessary to un- and re-install IntelliWatch. Naturally, you don't want to lose your current configuration settings in the process.

The following procedure assumes that you want to do all uninstalls, reinstalls, and updates on a system-by-system basis. If that's not the case, please adjust the following steps accordingly.

1 Before starting the rollout of the Domino incremental upgrade, designate one system as the repository of your IntelliWatch configuration settings.

All IntelliWatch product components should be installed on this system, to facilitate later replication.

- If you have a test system that can replicate with your production environment, simply install all IntelliWatch components on it, and replicate the databases to this system, rather than installing them from the media. If no servers in your test environment are able to replicate with your production systems, we suggest you select your least busy production system for carrying out Step 1, above.
- Since the setup allows you to switch to a different server during replication, should a given database not be found on a particular replication source server, you need not currently have a system in your environment where all IntelliWatch components are installed.
- Of course, if a particular component is not currently installed anywhere in your environment—Analyzer, for instance—there will be no settings to save, so simply elect to install any such databases from the media.

2 Uninstall IntelliWatch on a system before running the Domino incremental installer.

3 Run the Domino incremental installer.

4 Reinstall IntelliWatch—and replicate the IntelliWatch databases from the system discussed under Step 1, above.

2.1.4.2 Partitioned Notes servers

The IntelliWatch setup must be run once per Notes partition. Different IntelliWatch components may be installed on individual partitions. There is no limit on the number of partitions on a given system.

2.1.4.3 Pinnacle Console in a browser

No files are actually installed on the client machine. However, applets *are* downloaded from the server to which the Console is connecting. These applets are:

- Javasoft Swing package
- Pinnacle Console package

The first time a given workstation loads the Pinnacle Console, dialogs are displayed requesting acceptance of the IBM 'signature' on the packages in question.

The signature process enables applets to acquire the increased privileges necessary for full Pinnacle Console functionality. If you do not accept the IBM signature on *both* of these dialogs, the Console cannot function properly and does not load.

To avoid seeing these dialogs each time you connect to the Pinnacle Console, check the box at the lower left. Bear in mind that you are thereby accepting *all products bearing the IBM signature*.

2.1.4.4 IntelliWatch Stand-alone client

To install the stand-alone client, follow the above steps for the server installation, then connect to the browser version of the Pinnacle Console on that system. Click on the link to download the stand-alone client from the target server. (For instructions on

IntelliWatch connecting to the server, see *"Stand-alone IntelliWatch client" on page 46*.)

2.2.0.0 MA DEPLOYMENT

All IntelliWatch Management Agents in a given environment normally have the same Replica ID. This is accomplished by installing Management Agents from the media ***on the first server only***. All other servers obtain their Management Agents by Notes replication from the Primary Server.

This same procedure is used to propagate changes in MA configuration to all servers in your Notes environment.

2.3.0.0 SILENT SETUP

The Silent Setup was created to facilitate subsequent installations:

- on remote servers
- during scheduled down times

2.3.1.0 How does it work?

When the "Record Silent Setup Script" option is selected at the start of the setup (see *Figure 2-1*, below) the choices and entries you make are recorded in a script file (iwsetup.ini).

At the same time, a batch file is created (pinnacle.bat).

Subsequent installations of IntelliWatch can be carried out by running the batch file (which, in turn, calls the script).

2.3.1.1 Before creating the response file

Take a few minutes to:

- decide which components to install
- determine a useful name for the configuration to be recorded in the script

TO CREATE A RESPONSE FILE:

1 Run the IntelliWatch setup, and select the "Record a Silent Setup Script" option (see the arrow in *Figure 2-1*, below).

FIGURE 2-1: Select Silent Setup option

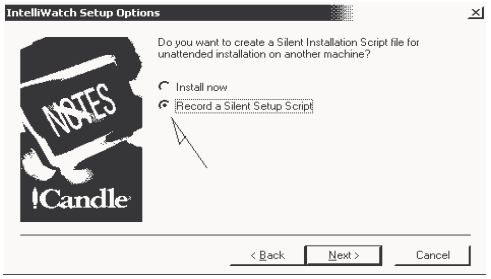

2 Assign a location and filename for the response file to be created.

This name and location are written to the batch file, which is placed with the response file in the specified location.

3 Step through the remaining setup dialogs, making entries appropriate to the desired IntelliWatch configuration.

2.3.1.2 Silent Setup on a new server

Once you've created the response file for a given IntelliWatch configuration, you're ready to carry out a Silent Setup on other servers (with that same configuration).

Two configuration files are required:

- response file
- batch file to run the Silent Setup

In addition, you need access to the IntelliWatch setup files.

TO RUN THE SILENT SETUP:

1. Place the two configuration files listed above in the directory you specified when creating the response file.
2. Launch the batch file, which includes the following contents:

```
<full path for setup executable>
  -b<full path for response file> -s
  -f1<full path for response file>
  -f2<full path for log file>
```

When you create the response and batch files, the above example syntax is on its own REM'd-out line. The following line of the file has the actual paths you specified/accepted when you ran the Silent Setup.

These values need only be changed if the file locations have changed.

Only -f2 is optional; all other flags are required. Flags must be used in the order shown.

The flags have the following functions:

- -b
  - IntelliWatch flag that identifies the location of the response file.
- -s
  - An InstallShield flag that causes the Silent Setup to run instead of the Standard Setup.
- -f1
  - An InstallShield flag that identifies the location of the response file.
- -f2
  - An InstallShield flag that records a successful setup, or reports an error code in case of failure.

*Do not confuse this file with **iw.log**, which details the steps taken by the Setup (on NT systems).*

2.3.1.3 Editing the response file

The Silent Setup script can be edited to change options for different servers.

Use any text editor to make changes. You must save the file in ASCII format, however, or the installation will fail.

2.3.1.4 What Silent Setup *won't* do

There is no opportunity to make changes in a given installation's configuration while the Silent Setup is running on a system.

2.4.0.0 INSTALLING ON WINDOWS 2000

2.4.1.0 Background

NT 4.0 offered various levels of registry and file-system security, but most of these options were ***off*** by default.

On Windows 2000, not only was the range of available security options increased, those options were turned ***on*** by default.

2.4.2.0 Potential issues

Not taking Win2K permissions into account when installing IntelliWatch can cause issues such as—but not limited to:

- Inadequate registry access for IntelliWatch components, which renders those components unable to get required configuration settings (and thus to run).
- Inadequate access to IntelliWatch executables (and/or the folders in which they reside), which prevents IntelliWatch components either from launching at all, or, if they launch, from running properly.

2.4.3.0 Avoiding issues

To avoid issues, install IntelliWatch using the following access level/login authority.

2.4.3.1 Minimum Access level

- Read-write access to the registry key:

(HKEY_LOCAL_MACHINE\SOFTWARE\Candle\IntelliWatch\[etc.])

- ·Read-write-execute access to the following files/folders:
  - Candle folder (and all subfolders)
  - IntelliWatch executables in the Domino executable directory
  - IntelliWatch databases (default location …/Data/IntelliWatch)

2.4.3.2 Login authority

The Login used to install IntelliWatch should have Full Control of all the above resources.

We recommend using an ID that has Administrator privileges while logged on to a Windows Domain—most often the account used to install Domino. Note that the Local Administrator lacks sufficient authority.

2.4.4.0 Sending IBM Tivoli Monitoring event messages from IBM Tivoli IntelliWatch (itwsditm utility)

To send IBM Tivoli Monitoring event messages from IntelliWatch on a Lotus Domino Server, IntelliWatch version 27.39 or later must be installed on the Domino Server.

The iwsditm IntelliWatch program sends IBM Tivoli Monitoring event messages to the agent that is monitoring the Domino Server. Like other IntelliWatch utilities, the iwsditm program can be run from the command line, a trigger, or a batch file to be called by IntelliWatch Advanced ServerWatch or IntelliWatch Crash Detection.

2.4.4.1 Command line

To run the iwsditm program from the command line, use the following command:

```
iwsditm /M:"messsage string"
  {/S:severity] [/H:hostname] [/P:port]
```

Where:

/S:severity - Priority of the message, optional: 1 = information (default), 2 = warning, 3 = alarm.

/H:hostname - Host name for the IBM Tivoli Monitoring agent, optional

/P:port - Port for the IBM Tivoli Monitoring agent, optional. The port is the same port specified in configuring the agent.

Example:

iwsditm /M:"Server Not Responding"

Trigger

When configuring a trigger, select the **Send ITM Event** command from the command list. When the trigger fires, this command sends the designated message to the Monitoring Agent for Lotus Domino.

1 From the IBM Tivoli IntelliWatch console, Monitor section, "Command" tub, select the "Create New Command" tool bar.
2 Select **Program** from the list.
3 In the **Name** field, type `Send ITM Event`.
4 In the **Program** section, type `iwsditm`.
5 In the **Parameter** field, type `/M:`*`MESSAGE`*.
6 Save the command.

Batch file

Use the following procedures to run the iwsditm program from a batch file that can be called by Advanced ServerWatch or Crash Detection.

Called from Advanced ServerWatch

To use the iwsditm program from a batch file that is called from Advanced ServerWatch, perform the following steps:

1 Create a batch file that contains the iwsditm command and uses %1 to pass the Advanced ServerWatch message. For example:

```
iwsditm /M:%1
```

2 From the IntelliWatch console Advanced ServerWatch tab, open or create an action profile by clicking **Hub > Profiles > Actions** to open the Modify Action Profile window.
3 In the Notification section, select the **Start Program** check box.
4 In the **Program to Run** field, type the location of the batch file.
5 Click **OK**.

Called from Crash Detection

To use the iwsditm program from a batch file that is called from Crash Detection, perform the following steps:

1 Create a batch file as described in "Called from Advanced ServerWatch".
2 From the IntelliWatch console Parameter Configuration tab, click Monitor > Crash Detection.
3 In the Value column, enter the information, as follows:

Name	Value
Start Program(1)	1
Start Program(2)	1
Program(1)	*Batch file location*, for example: C:\Tivoli\IntelliWatch]\COMMON\senditm.bat
Program(2)	*Batch file location*, for example: C:\Tivoli\IntelliWatch\COMMON\senditm.bat

For Start Program(1) and Start Program(2), the value of 1 enables the batch file program.

2.4.4.2 Integrating IntelliWatch with IBM Tivoli Monitoring 6.x

To integrate IntelliWatch with IBM Tivoli Monitoring, version 6.x, you need to understand how integration works. The Tivoli Monitoring Agent for Messaging and Collaboration retrieves statistics from a

Domino server as native Domino statistics as well as IntelliWatch statistics.

To enable the Tivoli Monitoring Agent for Messaging and Collaboration to retrieve these statistics, the port name that IntelliWatch uses to communicate with this monitoring agent must be defined. This is accomplished by setting up an instance for a Domino server using the IBM Tivoli Monitoring **Manage Services** option.

Note: You do not need to perform this step if the common name of the Domino server is identical to the hostname of the Domino server.

You also need to define the method for sending notifications from IntelliWatch to the Tivoli Monitoring Agent for Messaging and Collaboration. You do this by declaring the port that IntelliWatch uses to send messages to and from the Tivoli Monitoring Agent for Messaging and Collaboration.

Use the parameter database and insert the parameters, as follows:

Parameter for <domino server hostname>

Basics:

- Description: The hostname for the machine the monitoring agent is running on to send IntelliWatch alerts
- Server/Group name: <domino server hostname>

Set/Modify Parameters:

- Parameter section: Monitor
- Parameter category: USER
- Parameter to set: TCPIP Hostname
- Parameter value: <domino server hostname>

For all servers, you need to define the port that IntelliWatch uses to communicate with the Tivoli Monitoring Agent for Messaging and Collaboration. This is done by setting the following parameter in the parameter database.

Parameter for <your server list>

Basics:

- Description: The port IntelliWatch uses to communicate with the Tivoli Monitoring Agent for Messaging and Collaboration.
- Server/Group name: <your server list>

Set/Modify Parameters:

- Parameter section: Monitor
- Parameter category: USER
- Parameter to set: TCPIP Port
- Parameter value: 44576

The recommended port number is 44576. However, any port defined for the Tivoli Monitoring Agent for Messaging and Collaboration instance can be used by IntelliWatch.

For details on how the Domino agent works, refer to Monitoring for Messaging and Collaboration > Lotus Domino Agent User's Guide at http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?toc=/com.ibm.itm.doc/toc.xml

Chapter

Management Agents

3

IntelliWatch Monitor uses Notes databases called Management Agents (MAs) to provide local problem detection and correction for your Domino servers.

Triggers periodically check the local system for user-specified conditions, then execute configured Commands when the condition is detected.

Chapter Contents

3.1.0.0 OVERVIEW

Understanding IntelliWatch Monitor's design enables you to use the product more effectively.

Readers unfamiliar with Monitor fundamentals are encouraged to read the next several sections carefully. Experienced users may find that a quick scan of the material is sufficient.

3.1.1.0 MA Organization

Commands are housed in a single database (**iwcmd.nsf**), and Triggers are distributed across seven MAs, grouped by area of Domino functionality:

- IW Advanced Server (Notes Clustering)

  iwadvsrv.nsf
- IW Core Server

  iwserver.nsf
- IW Database Corruption

  iwcorrpt.nsf
- IW Internet Services

  iwnet.nsf
- IW Mail, C&S (Calendar/Schedule Connector)

  iwmscc.nsf
- IW Replication

  iwreplca.nsf
- IW SMTP MTA

  iwsmtp.nsf

3.1.2.0 Local monitoring

What is meant by *Local Monitoring*?

Simply put, IntelliWatch Monitor must be installed on all servers on which Triggers are to run. Even if a Trigger is configured to check for a condition on a particular server, it cannot do so if the product is not installed on that system.

3.1.2.1 What is a Trigger?

Triggers are Notes documents that contain four main sections:

- description

  – Information as to the Trigger's purpose and function in your environment.
- monitoring information

  – Parameters such as date, time, and custom monitoring frequency.
- conditional statements

  – Specifies those conditions under which you want the Trigger to fire.
- action lists

  – Specifies the actions to be taken by Monitor when Triggers evaluate to 'True'.

  (Action categories: Notification and Corrective.)

Out of the box, Monitor includes Triggers covering a broad range of Notes problems. Also included are templates that enable you to create custom Triggers to detect/correct issues peculiar to your environment. For details on individual Trigger types, see *pages 74 to 117*.

3.1.2.2 What is a Command?

Commands are Notes documents that contain two main sections:

- configurable parameters

Parameter details vary greatly, depending on the Command. For details on individual Command types, see *pages 118 to 153*.

- program or batch file
  - These can be standard Notes tasks (Router, HTTP, for example), as well as programs/batch files of your own choosing/creation.

3.1.3.0 MA Deployment

As discussed in the Installation Guide (Chapter 2), all IntelliWatch Management Agents in a given environment normally have the same Replica ID. This is accomplished by installing Management Agents from the media on the first server only. All other servers obtain their Management Agents by Notes replication from the Primary Server.

This same procedure is used to propagate changes in MA configuration to all servers in your Notes environment.

Accessing MAs on the Primary Server is therefore the first step in creating/editing Triggers and Commands.

3.1.4.0 Accessing MAs

IntelliWatch databases can be accessed in three ways:

- Pinnacle Console that runs in Internet Explorer
- stand-alone IntelliWatch client that installs on your Admin workstation
- directly via a Notes client

The first two access methods are outlined here. To access MAs via a Notes client, see *"Management Agents" on page 326*.

*With the exception of **console.nsf**, configuring IntelliWatch by accessing databases directly via a browser is not supported.*

3.1.4.1 Pinnacle Console: via a browser

Perhaps the biggest advantage of this method is that you need not be at your Admin workstation, but can use any NT machine in your environment that has Internet Explorer installed on it.

(Presupposes that you have the necessary access rights to IntelliWatch databases on the remote system, and that the Internet password option is properly configured.)

TO ACCESS MAS VIA A BROWSER:

1. Make sure that 1) the Domino server whose MAs you want to access is running, and 2) the HTTP task is loaded.
2. Bring up Internet Explorer.
3. Type the URL in the IE Address field, based on the following model:

```
http://[DomainNameOfPrimaryServer]/[IntelliWatchDataDir]/console.nsf
```

Depending on the configuration of your domain (and from where the Primary Server is being accessed), you may need to use the complete Internet domain name.

4 Press Enter.

5 Click on the link on the upper left of the page to launch the Console.

Depending on how your environment is configured—and whether you are accessing the Console for the first time—you may need to enter a User name and a password. For details, see *1.4.0 on page 24*.

6 Go to **Solutions > Management Agent** via the drop-down menus; alternatively, click on the toolbar icon.

7 Click on the tab of the desired MA to open it.

3.1.4.2 Security Mechanism

For the security mechanism governing access to the Console, see *1.4.0 on page 24*.

3.1.4.3 Stand-alone IntelliWatch client

You may prefer running the Pinnacle Console as a local, stand-alone application. (The Replication Check console utility is available *only* through the stand-alone client).

To install this version of the console, see the *"Installation Guide" on page 31*.

The steps required to launch the stand-alone client differ from those required to bring up the Console via Internet Explorer. Once the Console is launched, however, follow the same procedures for creating Triggers, editing them, and so on.

The first time you connect to the stand-alone client, you will be asked to provide the name of the Primary Server.

*This operation causes the **pinnacle.ini** file to be created and stored in the Candle/IntelliWatch/Console folder.*

The Primary Server is listed in this file, and subsequent connections will default to that server.

TO ACCESS MAS VIA THE STAND-ALONE CLIENT:

1 Go to **Start > Programs > IntelliWatch**, and click on the icon.

2 Go to **Solutions> Management Agent** via the drop-down menus; alternatively, click on the toolbar icon.

If so desired, you can configure the Pinnacle Console, when launched, to load the interface for a particular Solution.

*Just go to **File > Select Initial Solution** via the drop-down menus. This launches a list box containing all IntelliWatch Solutions. Select the one you want to have presented at start-up, and click OK.*

The interface for the selected Solution displays the next time you launch the Pinnacle Console.

3 Click on the tab of the desired MA to open it.

3.1.4.4 Accessing MAs via Notes

You may prefer accessing Management Agents via a Notes client.

The steps required to access MAs via Notes are the same as for any other database.

TO ACCESS MAS VIA NOTES:

1 Go to **File > Database > Open >** [**name of server**] > [**name of database**] via the drop-down menus.

2 Click Open.

All necessary access rights must be in place for all databases so accessed. At a minimum, appropriate permissions are required for the MA in question and the IntelliWatch Command database (**iwcmd.nsf**).

3.1.4.5 Accessing MAs on a different Server

Should you want to access MAs on a server other than your Primary Server, how you proceed depends on whether you are running the Pinnacle Console in a browser, or are using the stand-alone client.

TO CONNECT VIA A BROWSER:

1 Adjust the URL used to access **console.nsf**.

If the alternate server is in the same domain, and has the same directory structure as the Primary Server, simply change that part of the URL corresponding to the Domain Name of the server.

For example, if you were loading the Console using the URL http://PrimaryServer/IntelliWatch/console.nsf, adapt it to read http://OtherServer/IntelliWatch/console.nsf.

If you access the alternate server regularly, create a Favorite link to simplify the process.

2 Once you connect to the other server, load the Pinnacle Console by clicking on the link at the upper left of the page.

TO CONNECT VIA THE STAND-ALONE CLIENT:

1 Load the Console.

2 Go to **File > Primary Server** via the drop-down menus.

This brings up a Select Primary Server pop-up dialog.

3 Enter the name of the alternate server.

If the alternate server's Host Name differs from its Server Name, click on the Advanced button to enter the Host Name.

4 Click **OK** to connect.

3.1.4.6 Default connection procedure

The rest of this User's Guide assumes you will be carrying out the procedures described via the Pinnacle Console, rather than via Notes. Only in special situations are alternative procedures discussed.

Please refer to the preceding sections on alternative configuration methods when carrying out the procedures outlined in the rest of the manual. (See also *"Configuring IntelliWatch via Notes" on page 325*.)

3.1.5.0 MA Interface

Whether you access the Pinnacle Console via Internet Explorer, or are using the stand-alone client, the MA interface is identical.

3.1.5.1 Trigger databases

Seven MAs are created by default, each subdivided into the following six (tabbed) sections:

- Basic Information
- Monitoring Information
- Condition
- First Occurrence
- Second Occurrence
- Additional Occurrences

All Trigger sections but one—Condition—are the same for all Trigger types.

For field-by-field explanations of common Trigger sections, see *pages 65 to 73*.

For a discussion of Trigger fields by type, *pages 74 to 117*.

3.1.5.2 Command database

One Command database is created, serving all Management Agents. It, too, is accessed by a separate tab of the interface.

For a discussion of Command types, *pages 118 to 153*.

3.1.6.0 Trigger Mechanics

The Trigger process can be broken down into three stages (see graphic on following page):

- **Evaluation**
- **Action Lists**
- **Idle State**

A number of configurable options have been left out of the flow chart, below. For a step-by-step outline of Trigger functionality, see Figure 3-1 on page 53.

3.1.6.1 Evaluation

Monitor checks for a specified condition or value on servers for which a Trigger is enabled. Examples include: a file of a certain size; Mail.Dead exceeding a certain number; a particular ServerTask that's not running.

Evaluation has two possible outcomes:

- condition is not met

No *Action Lists* are executed. The Trigger goes into an *Idle State* until the start of the next monitoring cycle.

- condition is met

Action Lists are executed to notify/remedy, and the Trigger re-evaluates to determine if the monitored condition/value still exists.

For a flow chart depicting all possible evaluation paths for the three Action Lists, see *Figure 3-1* and section *3.1.6.3*, below.

3.1.6.2 Monitoring frequency

The monitoring frequency for IntelliWatch Triggers can be customized in two ways:

- setting for all Triggers
  - During the Setup, you are given the opportunity to accept the default value of 15 minutes, or to set a custom interval that applies to all Triggers. This value is stored in:
    - the registry (NT)

```
HKEY_LOCAL_MACHINE/SOFTWARE/Candle/IntelliWatch/Monitor/IW Agent/Monitoring Frequency
```

    - *iwmon.ini* (UNIX)

```
[IWAgent]
Monitoring Frequency
```

  The easiest way to change this value globally is by means of the Parameter Configuration Utility (see *"Parameter Configuration Utility" on page 350*).

- setting for individual Triggers
  - In the *Monitoring Frequency* field of a Trigger's Monitoring Information tab, enter a value *that will apply to this Trigger only.*

  If and when this MA is replicated to other servers, all copies of the Trigger will have the same custom monitoring frequency.

  Please note that custom monitoring frequencies are stored in the Trigger document only, and are not found in a global configuration repository.

While you may want the majority of enabled Triggers to evaluate at the 'global frequency', certain Triggers are most effective when configured to evaluate less (or more) frequently.

For example: Except in unusual circumstances (best known to local Admins), there should be no need to monitor Disk.C.Free every fifteen minutes.

Similarly—and subject to the dictates of local replication schedules—there should be no need to evaluate Replication Readiness or Replication Integrity Triggers as often as the global monitoring frequency.

3.1.6.3 Action Lists

Action Lists are IntelliWatch Commands you select, to be executed when a specified condition or value is detected.

Three Action Lists are available, each offering a range of options. Monitor's multiple Action Lists make possible escalated problem detection/correction.

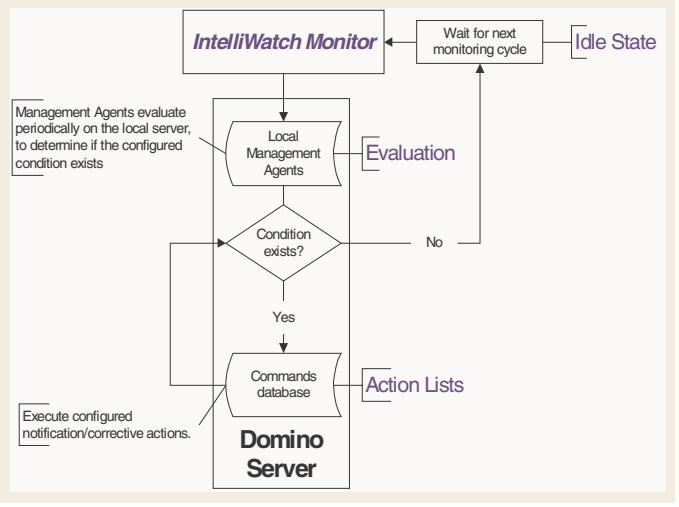

Available Commands range from e-mail notification of the Admin to Start Program Commands that run ServerTasks (COMPACT, for example, on databases with excessive White Space).

You can also create custom Commands (for more information, see *"Creating Commands" on page 58*).

Example 1: Initially, notify only

You're dealing with a condition that, by itself, isn't serious enough to warrant recycling the Notes server—unless it persists. Applying the principle of escalated problem detection/correction, you might proceed as follows:

- First Occurrence
  - notification actions only
- Second Occurrence
  - page Admin, and possibly stop and restart the server task in question
- Additional Occurrences
  - recycle the Notes server, or perhaps reboot the OS

    For details on configuring the number of Additional Occurrences, see *"Repeat:" on page 70*.

3.1.6.4 More on Action Lists

The escalated problem detection/correction afforded by Monitor's multiple Action Lists is not always required (or even relevant), however. For certain Trigger types, only the First Occurrence is relevant, as in the following two examples:

- latest Notes.rip Trigger
  - The Admin has configured a File Trigger to check for a Notes.rip with a date more recent than the *Latest Rip Date* value in the Pinnacle configuration repository:
    - the registry (NT)

```
HKEY_LOCAL_MACHINE/SOFTWARE/Candle/Inte
  lliWatch/Monitor/USER/Latest RIP Date
```

    - *iwmon.ini* (UNIX)

```
[USER]
Latest RIP Date
```

    As is normally the case, the Trigger is configured to update the *Date/Time or Bytes* field (the *Field* is set to *Latest Rip Date*). One of the selected First Occurrence notification Commands is e-mail notification of the Admin, with the Rip file as an attachment.

    Since the value of *Latest Rip Date* is updated by a Trigger so configured, the Second (and Additional) Occurrences are irrelevant.

    A Trigger thus configured resets, and fires again only when it detects a Rip file with a date later than the one listed in the systems' configuration repository.
- Replication Readiness Trigger
  - The Admin has configured a Replication Readiness Trigger to check databases prior to replication for potential issues that could cause the process to fail. Two outcomes are possible, both of which make the Second (and Additional) Occurrence Action Lists irrelevant;
    - If none of the checked conditions exists, the Trigger doesn't fire, and scheduled replication occurs normally. In other words, there was no First Occurrence (which is still to come).
    - If, on the other hand, one or more of the checked conditions is detected, the Trigger

fires, and, one way or another, changes to the database will be the result.

Corrective actions, whether initiated automatically by Monitor, or manually by the Admin, will presumably have resolved all potential replication issues. Again, the next time the Trigger evaluates, it will be looking for a First Occurrence of potential replication issues of a database in a new state.

In other cases, execution of Commands (notification, and perhaps others) is desired for *all* detected occurrences of an issue until it is resolved. To assure that this occurs, do the following:

- select Commands to be executed on the First, Second, and Additional Occurrences Action Lists
- enter a value of -1 in the Repeat field of the Additional Occurrences tab

Monitor now executes Commands every time the configured issue is detected, until it is resolved (at which time the Trigger resets, and again looks for a First Occurrence).

If a (positive) integer—not -1—is entered in the Repeat field, a condition could exist where the Trigger 'runs out of occurrences' before the issue is resolved (see also "Repeat:" on page 70).

Once Monitor 'runs out of occurrences', it continues to evaluate once per monitoring cycle—*but takes no actions*.

A Trigger thus configured resets (and the First Occurrence Action List again comes into play) *only* when Monitor detects that the condition no longer exists.

Put another way: the condition that caused the Trigger to fire in the first place (condition=True) must be resolved (condition=False) before the Trigger will reset.

Example 2: User Mail files too large

The Admin configures a Trigger to check the size of Users' mail files, and notification actions are selected (for both the Admin and the User concerned).

If only the First Occurrence Action List has Commands selected for execution, unless and until the database's size again falls below the configured threshold, the Trigger will not fire again—even though the problem has not been corrected. (The Trigger *evaluates*, but takes no actions, since none are configured.)

If Second and Additional Occurrences are configured, further (selected) notification actions will be taken.

Remember ...

The only way to guarantee that notifications are sent for as long as the problem persists is to 1) select Actions on all Occurrence tabs, and 2) set the Repeat field to -1 on the Additional Occurrences tab.

3.1.6.5 Idle State

The *Idle State* is the period during which Monitor is waiting for the next monitoring cycle. No evaluation is taking place, and no actions are executed.

Although the duration of the Idle State is based on the Monitoring Frequency specified during the Setup, individual Triggers can be configured to have their own Monitoring Frequency (which overrides the default), thus lengthening—or shortening, when desirable—the duration of the Idle State.

3.1.7.0 Customizing MAs

Should you want to use fewer Management Agents than the seven created by the Setup, either: 1) create your own MA(s), and import Triggers from the default databases; 2) consolidate MAs created by the Setup by copying Triggers.

Both of these operations must be carried out through the Notes client.

3.1.7.1 MA template

Observe the following when customizing Management Agents:

- MAs must be created using **iwmon.ntf** (at a Notes client)
  - This is the same template used by the Setup to create Management Agents.

This template should not be modified in any way, as this could cause the databases created to malfunction.

- MA's filename must be listed in the MA Database List parameter
  - Triggers are only evaluated if the filename of the MA containing them is included in the MA Database List parameter, found in the Monitor section of your system's IntelliWatch configuration:
    - the registry (NT)

```
HKEY_LOCAL_MACHINE/SOFTWARE/Candle/IntelliWatch/Monitor/IWAgent/MA Database List
```

    - *iwmon.ini* (UNIX)

```
[IWAgent]
MA Database List
```

  - Once created, custom MAs must be placed in the IntelliWatch data directory
  - This can be a custom folder, *but* all Management Agents must be in one folder. Having multiple MA locations on a given system is not supported.

3.1.7.2 Controlling which MAs are displayed

When the Management Agent Solution is accessed at the Pinnacle Console, interfaces are created only for those MAs listed in the MA Database List parameter.

3.2.0.0 COMMON TRIGGER FIELDS

See *on pages 64 to 72* for descriptions of common Trigger dialogs.

FIGURE 3-1: Flow chart of IntelliWatch Monitor Triggers

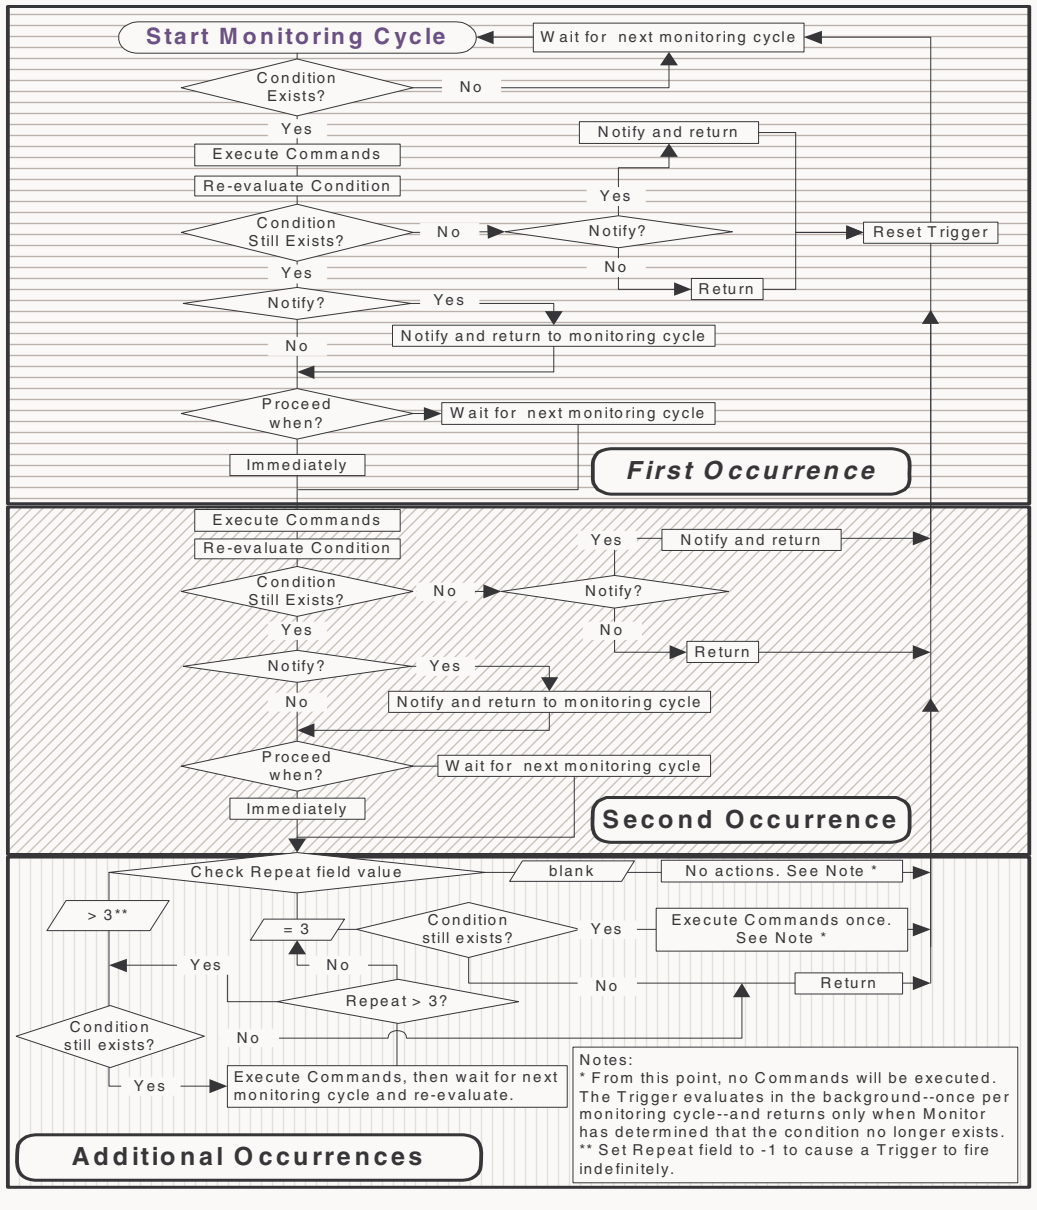

3.3.0.0 WORKING WITH TRIGGERS

Now that you've been introduced to overall Trigger functionality and design it's time to work with the documents themselves.

The following sections outline the procedures required for:

- **Creating Triggers**
- **Editing Triggers**
- **Enabling/Disabling Triggers**
- **Deleting Triggers**

The procedure outlined below presupposes you've already brought up the Management Agent Solution via the IntelliWatch UI.

If that is not the case, please follow the steps at "Accessing MAs" on page 45, then resume working with Triggers.

3.3.1.0 Procedures

3.3.1.1 Creating Triggers

The first step is deciding in which database you want to the Trigger to reside.

Remember ...

Base that decision on organizational convenience, since a Trigger's functionality is unaffected by the MA in which it resides.

In other words, you can create a Trigger governing database corruption in the Core Server MA, even though there is a separate Database Corruption MA.

As noted at *3.1.5.1*, Triggers consist of several tabbed sections, all but one of which—the Condition tab— is the same for all Trigger types.

Common Trigger sections are largely self-explanatory. Consult the page referred to in parentheses after each section, for usage information:

- Basic Information (*on page 64*)
- Monitoring Information (*on page 66*)
- First Occurrence (*on page 68*)
- Second Occurrence (*on page 68*)
- Additional Occurrences (*on page 70*)

The Condition section varies too widely from one Trigger type to another to make possible a common set of instructions. Usage information by type is arranged alphabetically, and included on *Trigger Type: on pages 74 to 116*.

Command options are discussed Occurrences.

Message options are discussed separately, *on page 72*.

3.3.1.2 Editing Triggers

Editing Triggers is usually a straightforward procedure, involving modifications to a few fields (or perhaps only a single parameter).

TO EDIT TRIGGERS:

1 Click on the IW MA tab containing the Trigger to be edited.

This displays a tree view of the Triggers in the selected IW MA.

2 If necessary, expand the Trigger list by dragging the viewing bar to the left or right to increase or decrease the MA viewing window.

3 Click the + block to the left of the Trigger type to expand the list.

4 Click the Trigger to be edited.

This displays the Trigger template listing its details.

5 Go to **File > Edit** via the drop-down menus; alternatively, click on the  toolbar icon.

6 Make the necessary changes to the Trigger field(s).

7 Go to **File > Save** via the drop-down menus; alternatively, click on the  toolbar icon.

3.3.1.3 Enabling/Disabling Triggers

The importance of managing which Triggers on your system are enabled/disabled should not be underestimated.

As with any application/process, Trigger evaluation and execution take up system resources. To make the most efficient use of Monitor Triggers, spend a few extra minutes when selecting the servers for which Triggers are enabled/disabled.

For example, if you have servers with ample free disk space, you may want to monitor it less frequently than you would on systems where disk space is limited.

TO ENABLE/DISABLE TRIGGERS:

1 From the Management Agent Console, click on the tab of the IW MA that contains the Trigger to be enabled/disabled.

Triggers are displayed in a tree structure in the list box to the left.

2 If necessary, adjust the size of the list box containing the Triggers by dragging the viewing bar to the left or right.

3 Click on the + block to the left of the Trigger type to expand the Trigger list.

4 Select the Trigger to be enabled/disabled.

5 To *Enable* Triggers, go to **File > Enable Triggers** via the drop-down menus; alternatively, click on the  toolbar icon.

6 To *Disable* Triggers, go to **File > Disable Triggers** via the drop-down menus; alternatively, click on the  toolbar icon.

3.3.1.4 Deleting Triggers

We suggest disabling Triggers rather than deleting them, *since there is no Undo for Trigger deletion*.

Should you inadvertently delete a Trigger, however, remember that only the open MA is affected. All other MAs that contained the Trigger, including the database template, still contain it. Subject to how you manage your MAs, you should be able to restore an inadvertently deleted Trigger via Notes replication.

TO DELETE TRIGGERS:

1 From the Management Agent Console, click on the tab of the IW MA that contains the Trigger to be deleted.

Triggers are displayed in a tree structure in the list box to the left.

2 If necessary, adjust the size of the list box containing the Triggers by dragging the viewing bar to the left or right.

3 Click the + block to the left of the Trigger type to expand the Trigger list.
4 Click the applicable Trigger to be deleted.

This displays the Trigger template.

5 Go to **File > Delete** via the drop-down menus; alternatively, click on the  toolbar icon.

A confirmation dialog is displayed.

6 Click Yes (the default) to confirm the deletion.

3.3.1.5 Copying Triggers

Copying Triggers is not possible using the Pinnacle Console. This operation must be carried out via a Notes client (see *10.2.1.4*).

3.3.2.0 Trigger Condition fields

Condition sections are discussed in alphabetical order by Trigger type (on *pages 74 to 117*). Information is organized as follows:

- **Basics...**
  - conditions that can be monitored
  - options available on selected operating systems
  - cautions
- **Putting it into practice...**
  - usage tips
  - real-world examples
- **Figure**
  - explanation of condition fields
  - pop-up option lists (where applicable)

For an individual Trigger type, see the alphabetically organized section *pages 74 to 116*.

3.3.2.1 The <VALUE> keyword

If you include this keyword in a message, the contents of the Trigger's conditional field are inserted in place of the keyword <VALUE>. The data represented by <VALUE> differ by Trigger type:

- Application Triggers
  - name of the application
- Database Triggers
  - contents of the field of the database
- File Triggers
  - filename
- Statistic Triggers
  - value of the statistic

3.3.3.0 Command fields by type

IntelliWatch Monitor Commands vary widely by type. While Commands such as IWSleep consist of nothing more than a user-specified Name and a Timeout, others contain an extensive list of configurable parameters.

Command types are presented in alphabetical order (on *pages 118 to 153*). Information is organized as follows:

- **Basics ...**
  - actions performed
- **Putting it into practice ...**
  - usage tips
- **Figure**
  - explanation of Command fields

For an individual Command type, see the alphabetically organized section *Command Type: on pages 118 to Command Type:152*.

3.3.4.0 Making use of keywords

IntelliWatch Monitor allows for the use of <KEYWORDS> in Triggers and Commands. These keywords simplify—and enhance—the power of Monitor.

Example 3:

You want to automate the removal of old documents from Notes databases.

Set up a Document Time-out Trigger to check for documents older than X minutes.

In the database field of the condition tab, enter *.nsf to search all databases. (Configure the other parameters in keeping with your environment.)

Configure the First Occurrence Actions List to include a Move/Remove Command that makes use of two IntelliWatch keywords:

- <DATABASE>
- <DOCUMENT_ID>

Now, whenever the Trigger detects a document older than the configured timeout, it will pass the name of the database to the Move/Remove Command, along with the Document ID. The Command will move the document to a new database, or delete it, depending your configuration choices.

3.4.0.0 WORKING WITH COMMANDS

Just as with Triggers, IntelliWatch Commands can be modified (or created) to suit the needs of your environment.

The following sections outline the procedures required for:

- **Editing Commands**
- **Creating Commands**
- **Deleting Commands**

The procedures outlined below presuppose you've already brought up the Management Agent Solution via the IntelliWatch UI. If that is not the case, please follow the steps at "Accessing MAs" on page 45, then resume working with Commands.

3.4.1.0 Procedures

3.4.1.1 Editing Commands

Editing Commands is usually a straightforward procedure, often involving modification of a single field.

TO EDIT COMMANDS:

1. From the Management Agent Console, click on the Commands tab to display the Commands database (**iwcmd.nsf**).
2. If necessary, adjust the size of the list box containing the Commands by dragging the viewing bar to the left or right.

3 Click the + to the left of the desired Command type to expand the list.
4 Click the Command to be edited.
5 Go to **File > Edit** via the drop-down menus; alternatively, click on the  toolbar icon.
6 Make the necessary changes to the Command fields.

For more details about these fields, refer to *"Command fields by type" on page 56*.

7 Go to **File > Save** via the drop-down menus; alternatively, click on the  toolbar icon.

3.4.1.2 Creating Commands

Before creating a new Command, give some thought to a Name that is *short* and *unique*. The Name field is used to populate the Available Commands list box; the more descriptive the Name, the easier it is to select Commands quickly and accurately.

TO CREATE COMMANDS:

1 Go to **File > New > Create Command** via the drop-down menus; alternatively, click on the  toolbar icon.

The cascading Command type menu is displayed.

2 From the cascading menu, click the Command type to be created.

This brings up the Command template appropriate to the selected type.

3 Fill in Command fields (including any available optional parameters you want to use).
4 Click Ok to save the new Command; alternatively, go to **File > Save** via the Management Agent drop-down menus.

3.4.1.3 Deleting Commands

We suggest deselecting Commands rather than deleting them. For two important reasons:

- Deleting a Command affects all local MAs whose Triggers invoke it.
- Deleted Commands cannot be restored.

Until you replicate **iwcmd.nsf**, a deleted Command still exists in copies of the database residing on other servers. Subject to how you manage your IntelliWatch databases, the Command can be restored via Notes replication.

TO DELETE COMMANDS:

1 From the Management Agent Console, click the Commands tab.
2 If necessary, adjust the size of the list box containing the Commands by dragging the viewing bar to the left or right.
3 Click the + to the left of the Command type to expand the list.
4 Select the Command to be deleted.
5 Go to **File > Delete** via the drop-down menus; alternatively, click on the  toolbar icon.

3.4.1.4 Copying Commands

Copying Commands is not possible using the Pinnacle Console. This operation must be carried out via a Notes client.

3.4.1.5 Location of Commands Database (iwcmd.nsf)

Commands are kept in a separate database (**iwcmd.nsf**) , located by default in the **/intelliwatch** subdirectory (of the Notes data directory). Should you want to move this

database, add the following variable to the **Notes.ini** *on the server:*

```
$IWCommandDb=<directory>\iwcmd.nsf
```

The <directory> should be relative to the Notes data directory.

If you store iwcmd.nsf in a location other than the default, but fail to add the above line to your Notes.ini, an error occurs. Working with Triggers cannot resume until this condition has been remedied.

3.5.0.0 MONITOR'S AVAILABILITY STATS

3.5.1.0 How they work

3.5.1.1 Up and Down documents

Monitor's Availability Reports are calculated on the basis of the Up and Down documents in the Server Status Log view of the iwstats.nsf database—if you are running maintenance release 27.36 or higher, your Up/Down documents may be stored in the iwstatus.nsf (see sections *3.5.1.3* and *3.5.1.6*, below).

Availability is calculated daily at a time specified in the NT registry, or the iwmon.ini on UNIX (see the Calculation Time parameter at the right of Figure *Figure 3-2*). (Each time IntelliWatch Monitor initializes, an Up document is created; each time Monitor shuts down, a corresponding Down document is created.)

3.5.1.2 How availability is calculated

A server is deemed Available between an Up document and the next Down document, and Unavailable between a Down document and the next Up document.

However, a server's first Up document in a given calculation period, is treated differently (it may be the Up document created the first time Monitor initializes after installation or after a purge of the database, e.g.).

Currently, Monitor's Availability Statistics do NOT make an allowance for planned unavailability—such as when the server is brought down for maintenance. (This feature may be added as en enhancement in a future maintenance release.)

3.5.1.3 Changes in 27.36

Starting with maintenance release 27.36, the option exists of using an alternative database to store the Server Status Log on which Monitor's Availability Reports are based, namely **iwstatus.nsf**.

*The name **iwstatus.nsf** is hardcoded; changing it is not supported.*

3.5.1.4 Why this new database?

Occasionally, customers have experienced errors in Monitor's Availability Statistics, due to a database purge interval that was shorter than their maintenance interval.

Consequently, Up and Down documents (in the Server Status Log View) got out of sync, resulting in errors.

3.5.1.5 May I continue to use iwstats.nsf?

Yes. No changes have been made in the design of **iwstats.nsf**, so you still have the option of using that database to store the Server Status Log.

FIGURE 3-2: As seen using the Parameter Configuration Utility of the Pinnacle Console

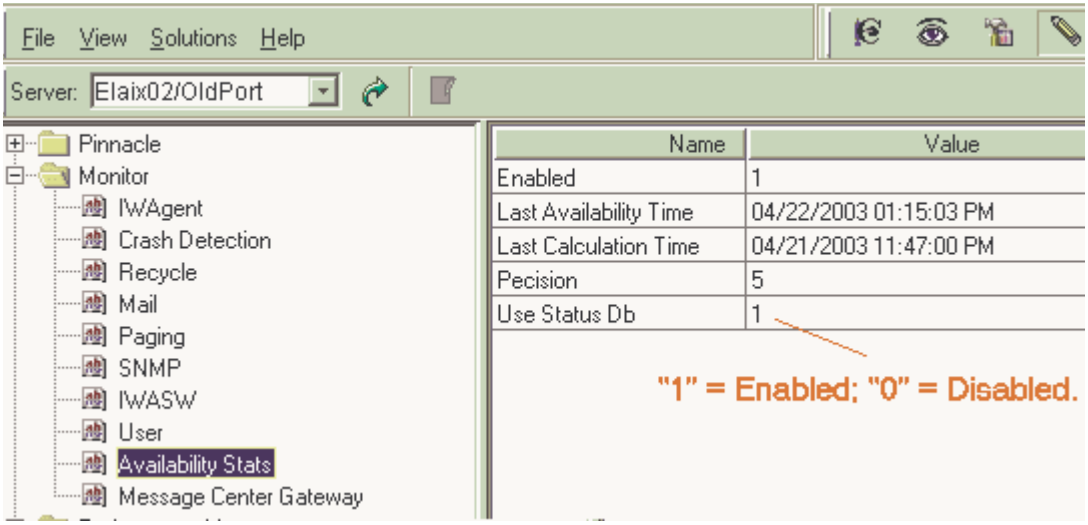

3.5.1.6 What's the default database for Availability Statistics?

That depends on whether you're doing a fresh install of 27.36, or updating an existing installation.

- Fresh installs

Fresh installs of 27.36 will have **iwstatus.nsf** as the default database.

- Upgrades (i.e. using Update####.exe)

When updating a pre-27.36 version of IntelliWatch, **iwstats.nsf** is retained as the default.

The new setting (Use Status Db) is created, but set to "0", i.e. disabled.

3.6.0.0 FAQs

Q: When I connect to a Primary Server, not all MA interfaces are created at the Pinnacle Console. Why is this?

A: Interfaces are not created on the basis of the MAs present on a machine, but as per the **MA Database List** value in the **IWAgent** section of your IntelliWatch configuration (on the server to which you're connecting).

Interfaces are created *only* for the MAs in this list.

Adjust the value in the MA Database List to include the MAs you want to view at the Console, then reconnect to the server. This value is located in:

– the registry (NT)

```
HKEY_LOCAL_MACHINE/SOFTWARE/Candle/Inte
  lliWatch/Monitor/IWAgent/MA Database
  List
```

– *iwmon.ini* (UNIX)

```
[IWAgent]

MA Database List
```

Iwagent also refers to the MA Database List value, to determine which Triggers should be evaluated. The Triggers in any MAs not in this list are not evaluated—even if they are enabled.

Q: Do I have to configure MAs on each server individually?

A: *By no means.* Just follow this simple procedure:

1 Install Monitor on your primary server. When given the option of installing MAs from the media or replicating from another server, choose the former.

2 Install Monitor on your other servers. Instead of creating the MAs from the installation media, however, replicate them from the primary server. All the MAs in your environment now have the same Replica ID.

3 Whenever you make changes to the MAs on your primary server, replicate them around your environment to maintain consistency.

Q: Our environment uses Domino (two-server) clusters, and File Triggers check the size of mail files. Since the replica copies on both servers of the cluster are the same size, and since the Triggers are enabled on both servers, why do I get messages from only one server?

A: Let's take the example of an environment where ServerA is clustered with Server B.

Although User1's home server is ServerA, a replica copy of her mail file exists on ServerB. However, when a File Trigger evaluates on ServerB's copy of the mail file, and checks the local copy of the Domino Directory for users associated with that server, User1 is deemed NOT FOUND.

With Pinnacle 99, unless the admin created a 'NOT FOUND' user on ServerB, the message ended up in Dead Mail. This functionality was adapted in IntelliWatch; when a mail user is NOT FOUND, no attempt is made to send messages.

Q: From time to time, the statistics contained in the Availability Report View of iwstats.nsf report 0% availability, even though the server has been running and IntelliWatch Monitor is otherwise functioning without issues.

If you're running maintenance release 27.36 or above, this issue can be avoided using new functionality added in that version of IntelliWatch (see the relevant sections 3.5.1.3 and 3.5.1.6, above).

Users running maintenance releases prior to 27.36 may resolve the issue by following the steps outlined below—although we recommend upgrading to 27.36 (followed by configuring your systems to use the new database for storing Up/Down documents).

A: Monitor's Availability Report is calculated on the basis of the Up and Down documents in another view in iwstats.nsf, namely the Server Status Log. (Each time IntelliWatch Monitor initializes, an Up document is created; each time Monitor shuts down, a corresponding Down document is created.)

The fundamental cause of this issue is the deletion of Up/Down documents from the Server Status Log View in iwstats.nsf. The most common cause of Up/Down document deletion is a Purge Interval for iwstats.nsf that is shorter than the server's maintenance interval.

The most straightforward remedy is to:

1. Delete all Up/Down documents on the LOCAL server (on each server affected by this issue).
2. Replicate iwstats.nsf to the system where statistics for all servers are collected to assure environment-wide consistency in this database.
3. Recycle Monitor on each affected server.
4. Verify that a Down and an Up document were created by Step 3.
5. Wait till the next calculation time (by default the evening of the day when Monitor was recycled, although, in some cases, an additional day has been required).

Starting with maintenance release 27.36, users will have the option of having these documents stored in a separate database, iwstatus.nsf. For details, see

Q: IntelliWatch Application Triggers are not correctly recognizing when add-in tasks are running, causing Triggers to fire (or not fire) improperly.

A: In most cases, this occurs for one reason:

- the "Notes Add-in Task" was not selected

The binaries for Notes add-in tasks running on Wintel systems have an "n" prepended to the name of the executable, and also to the name of the task (just as an "a" was prepended on NT Alpha systems).

The corresponding binaries (and add-in tasks) on UNIX use the 'generic' name, i.e. with no prepended characters.

To cite an example, "router" on UNIX is "nrouter" on Wintel.

Selecting the "Notes Add-in Task" checkbox means that:

- an "n" is prepended internally to the listed task name on Wintel systems, but
- no characters are prepended when the Trigger is running on UNIX systems.

For more on this issue, see IntelliWatch Technotes on ibm.com.

Q: You find that, over time, the iwstats.nsf database (the repository for IntelliWatch statistics) in your environment grows very large and is replicated throughout your environment. This results in considerable wasted disk space (and bandwidth comsumption) if left unchecked.

A: Configure your systems so that only one copy of iwstats.nsf contains data for all servers in your environment, namely the copy that resides on the Analyzer server. In short, we suggest that you configure iwstats.nsf for one-way replication.

There are two ways to set up this one-way replication:

- Replication formulas
- ACL restrictions

We recommend the second method, since it's easier to implement and maintain.

The ACL entry for the Hub server itself (that is, where IntelliWatch Analyzer runs) should assign that server READER access. The ACL entry for the Spoke servers (where the data is being collected) should assign those servers DEPOSITOR access, as in Figure 2Figure 2, below.

Even using the above scheme to control the growth of iwstats.nsf on your servers, eventually you'll need to warehouse your IntelliWatch data. Use the database's Space Saver settings to accomplish this easily and quickly.

1. Open iwstats.nsf.
2. Go to **File > Replication > Settings > Space Savers** to configure the Space Saver settings.

Three days should be adequate on the spokes; make the setting as long as needed on the Hub where Analyzer runs.

Normally the Hub will hold 45 days of data, and will have a program in place to make a copy of the database every 30 days, so that individual months of data can be recovered as required.

Q: Why doesn't an Availability Trigger fire, which is configured to check the TCPIP port for SMTP (port 25), since the SMTP task is not running?

A: The TCPIP/Port option for this Trigger type checks for port availability, NOT to see if the port is being used by the protocol in question.

Basic Information

Basics ...

Only one field on the Basic Information tab impacts Trigger functionality: *Enabled*. A Trigger is evaluated *only* if Enabled is checked (and the MA in question is listed in the MA Database List parameter of your Monitor configuration).

All other Basic Information fields *are for informational purposes only*.

Therefore, be sure Trigger Names accurately reflect the settings on the Trigger's Condition tab. Otherwise, you might enable a Trigger on the basis of the Name, only to discover that it was actually configured to do something else.

For example: You have a Trigger named "Search for the word "Error" in Notes.log", but (in the Value field of the Condition tab) it's configured to search for the string "Replication Conflict".

This Trigger fires (or not) based on the presence of the string "Replication Conflict"—*not the word "Error."*

Avoid redundancy in your use of the Name, Explanation, Comments and Category fields. (See *Usage tips*, below.)

Putting it into practice ...

Name:
The number of characters displayed in this field varies due to several factors. Therefore, write 'front-loaded' descriptions that enable you—at a glance—to tell one Trigger from another.

Explanation:
Text in this field can be appended to messages sent by Triggers (use the <EXPLANATION> keyword).

Comments:
Some predefined Triggers include information here on how to enhance the Trigger. Use this field to record configuration tips that you want to keep with a given Trigger, but that you don't want sent in Notifications.

Categories:
Use to organize Triggers for your convenience.

Enable:
Checked if a Trigger is enabled, otherwise unchecked. Enabled Triggers are activated immediately after the Trigger is saved.

IntelliWatch Monitor does not process disabled Triggers.

Trigger: Basic Information

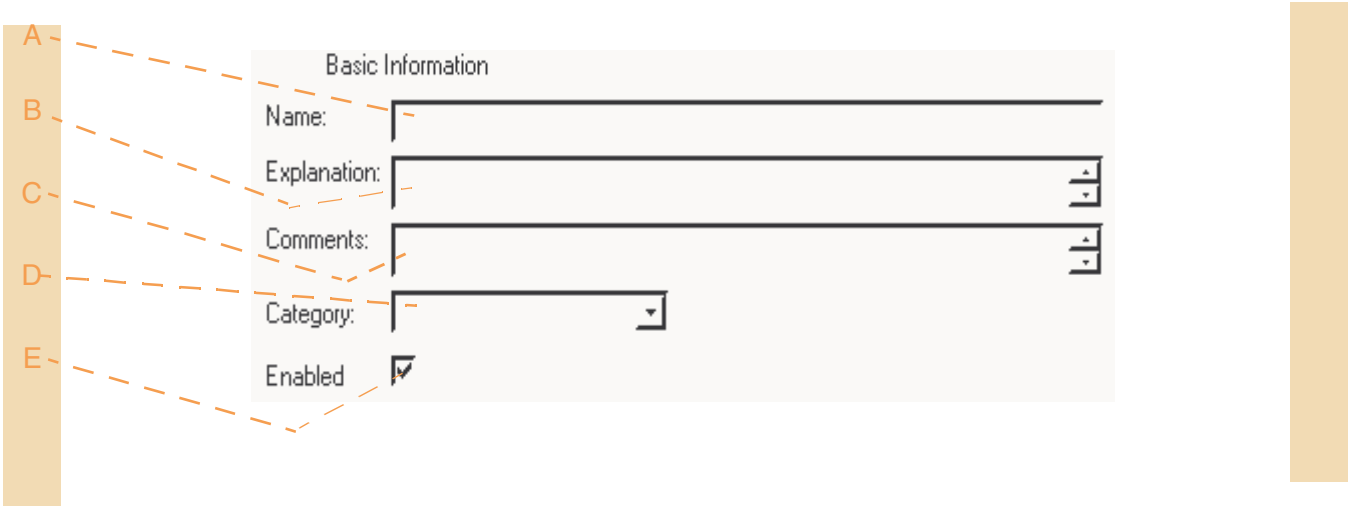

A: name of the Trigger
B: supplemental information as to what a Trigger does
C: configuration notes
D: Trigger category (double-click icon for examples)
E: checked if Enabled, otherwise unchecked

Monitoring Information

Basics ...

Provides information concerning:

- *where* (on which servers)
- *when* (days and times)
- how frequently to evaluate enabled Triggers

Another option available here is the *exclusion* of selected servers from being monitored by this Trigger. If this option has been checked, a Trigger is evaluated on *all* servers *except* those selected in the server list.

Triggers that should only be evaluated on R5 and above can be so designated on this tab.

When you select Statistic Generation for this Trigger, a count is incremented each time the Trigger fires.

All fields on the Monitoring Information tab are optional.

Putting it into practice ...

Monitoring Frequency:
Leave blank unless the frequency for this Trigger should deviate from the default.

Begin Time:
Time when Trigger starts to be evaluated. If blank, enabled at 12:00 AM each day. Valid times are from 12:00 AM through 11:59 PM. (Format: HH:MM.)

A 24-hour clock is assumed, unless you add AM or PM.

End Time:
Time when Trigger stops being evaluated. If blank, enabled till 11:59 PM each day. (Format as Begin Time.)

Monitor on Selected Days:
Select day(s) of the week on which to evaluate Trigger. If no boxes are checked, the Trigger evaluates every day.

Server List:
If the field is blank, the Trigger evaluates on all servers (unless the following option is selected).

Do not evaluate this Trigger...:
Excludes listed servers from Trigger evaluation.

Evaluate on Domino 5.0 and up:
Check to evaluate Trigger only on R5 and higher.

Generate Trigger statistic:
Creates statistic for each server on which the Trigger is enabled, counting how many times Trigger has fired.

Trigger statistics are in-memory only.

Statistic ID Field:
Part of the statistic name. Pick something that is both *unique* and *easily identifiable*.

Trigger: Monitoring Information

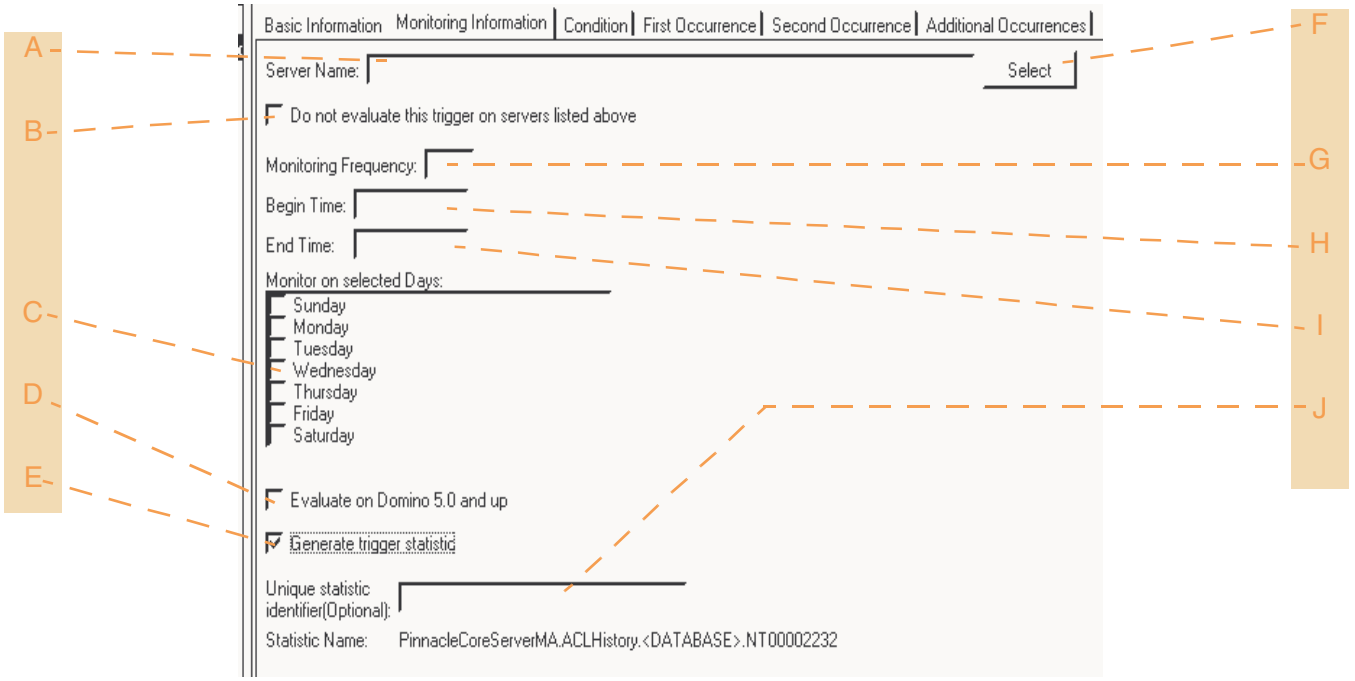

All of the following fields are optional.

- A: server list (group(s) or individual server(s))
- B: exclusion option
- C: day(s) on which to monitor
- D: if selected, evaluate Trigger only on R5 and above
- E: generate statistic every time Trigger fires
- F: button launches Select Server dialog
- G: monitoring frequency (if other than default)
- H: if field contains a value, not evaluated *before* this time
- I: if field contains a value, not evaluated *after* this time
- J: customization of name of generated statistic

First and Second Occurrences

Basics ...

The First and Second Occurrence tabs are identical, and so are discussed together. *However*, they must be configured individually.

The *First Occurrence Action List* consists of those IntelliWatch Commands that you want executed as soon as Monitor detects an issue on the server. In many cases, only notification actions are called for; in other situations, you'll want to configure Monitor to take immediate corrective actions (such as restarting an add-in task).

The *Second Occurrence Action List* comes into play only when the *First Occurrence Action List* did not successfully resolve the issue (or when a condition for which only notification actions were taken still exists). Therefore, the *Second Occurrence Action List* is usually more aggressive.

Putting it into practice ...

Available:

Highlight the Command(s) to be issued, then click on the Add button to move them to the Selected list box.

Use Shift and Ctrl keys in the usual way (under Windows) to select contiguous/non-contiguous Commands.

Selected:

Commands in the list appear *in order of execution.* Therefore, select the Commands in the order you want them executed, *not* in the order they appear in the list.

To change the order of a Command, highlight it, then use the Move Up and Move Down buttons.

To deselect Commands, first highlight them, then use the Remove button.

When to execute next action list:

Determines when to execute the next Action List, immediately or at the next monitoring cycle.

This option is not available for Database and ACL Triggers, but it can be extremely useful with all other Trigger types in accelerating multi-step problem resolution.

Trigger: First and Second Occurrences

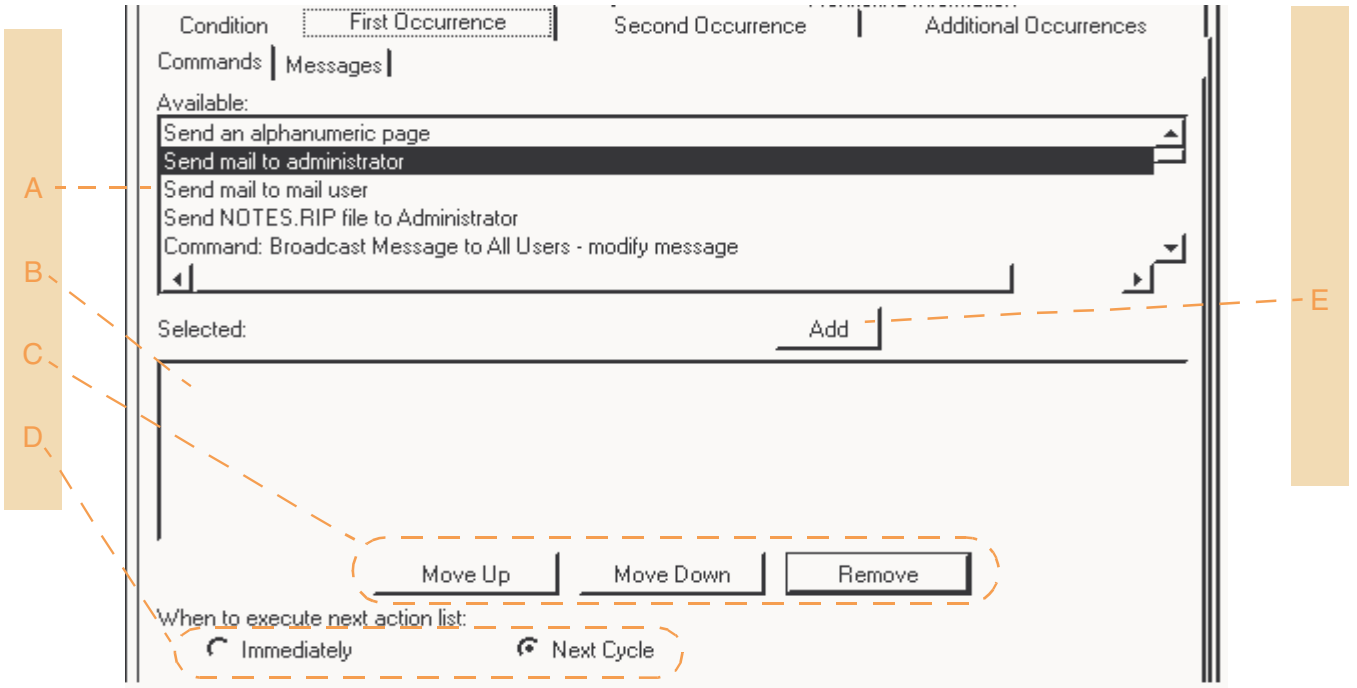

A: list of available Commands
B: list of selected Commands
C: navigation buttons
D: option of when to execute the next Action List
E: button for adding Available Commands to the Selected list

Additional Occurrences

Basics ...

The *Additional Occurrences Action List* is reached only if the first and second Action Lists have failed to resolve an issue (or if a condition still exists for which, until now, only notification Commands have been executed).

Usually, this is the most aggressive of the three Action Lists, although, depending on the issue, you may still choose to notify only.

Functionally, the *Additional Occurrences Action List* differs in two respects from the first two Action Lists:

- There is no option to execute the next set of Commands *immediately* or at the *next cycle*.
  - The *Additional Occurrences Action List* may be executed more than once, but successive executions occur *once per monitoring cycle*.
- Repeat field (see Usage tips, below)

Putting it into practice ...

Command selection procedures are identical to those for the first two occurrences. (See *"First and Second Occurrences" on page 68*.)

Repeat:

The Repeat field is unique to the *Additional Occurrences Action List*. The value entered represents *the occurrence(s) on which the Action List is executed*. To arrive at the *number of times* this Action List will be executed, simply subtract 2 (representing the first and second occurrences). To illustrate:

- A value of 3 means the Action List executes *one time* only: on the third occurrence (3 - 2 = 1). *Should the issue persist, this Action List is not executed on subsequent occurrences.*
- A value of 5 means the Action List executes on the third through the fifth occurrences, or a total of three times (5 - 2 = 3).

To cause the *Additional Occurrences Action List* to execute indefinitely (as long as the issue persists), enter a value of -1.

If configured Additional Occurrences is exceeded:

If a problem still exists after the configured number of Additional Occurrences has been executed, Monitor re-evaluates the condition once per monitoring cycle—*but takes no actions*.

The Trigger resets if and when:

- Monitor detects the issue has been resolved
- IWAGENT is stopped and restarted

Trigger: Additional Occurrences

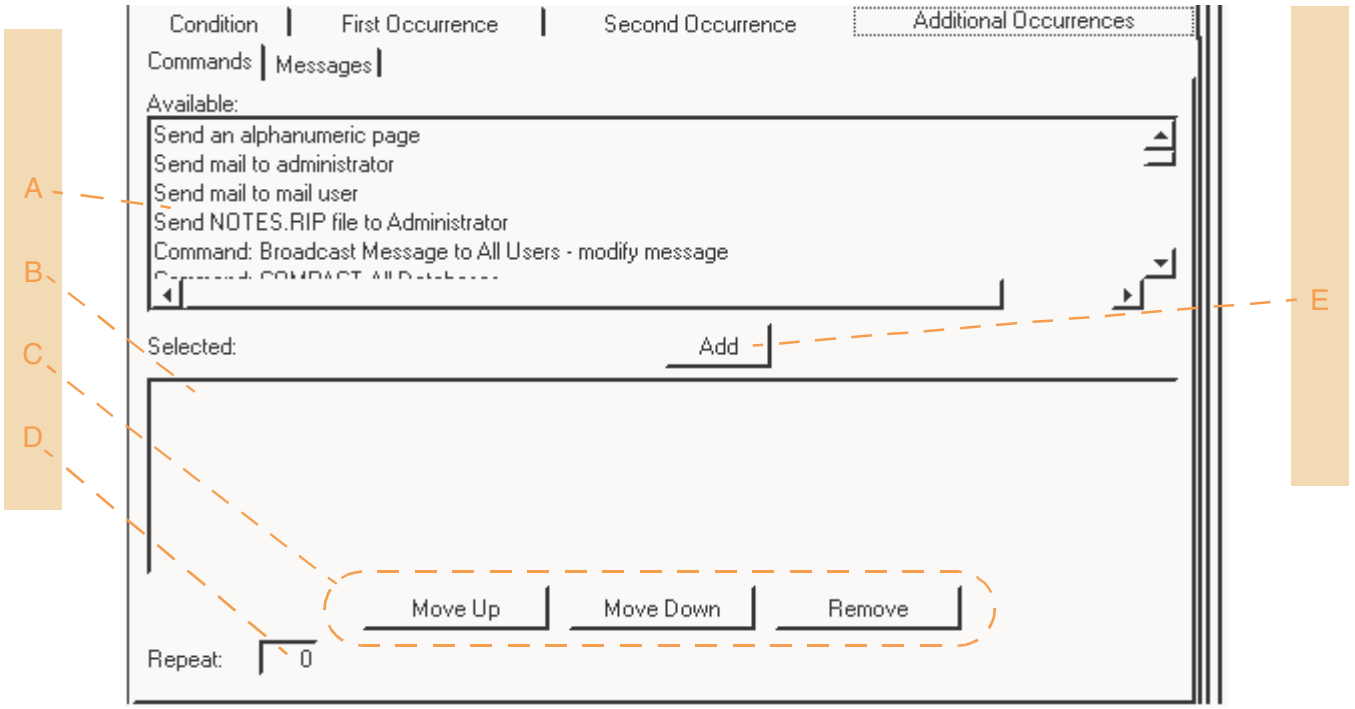

A: list of available Commands
B: list of selected Commands
C: navigation buttons
D: Repeat field
E: button for adding Available Commands to the Selected list

Messages

Basics ...

Allows you to tailor notification to your needs.

The three options are:

- when Commands are executed
- when Trigger re-evaluates
  - condition still exists
  - condition no longer exists

Database and ACL History Triggers have only the first of these fields.

Remember ...

At least the first of the Message fields has to be filled in, or the Trigger can't be saved.

The <VALUE> keyword:

For usage information, see *3.3.2.1 on page 56*.

Putting it into practice ...

Getting the most out of Trigger messages:

On the First and Second Occurrences, Triggers re-evaluate immediately after Commands are executed, to see if the monitored condition has been corrected.

Any active notification Commands for the Trigger in question send a message reporting on the condition.

This initial message, however, does not tell you if the executed Commands remedied the issue.

For that reason, most Trigger types include two additional messaging options:

- condition still exists
- condition no longer exists

Include an appropriate message in these latter fields, to be informed of the results of Trigger re-evaluation.

Why Database and ACL History Triggers lack these messaging options:

With these two Trigger types, no actions are taken to remedy the Triggering condition, making *Condition Still Exists* and *Condition No Longer Exists* inapplicable.

Database Triggers evaluate to True when a user-specified string has been detected, but no actions are taken to remove that string.

(Actions may well be taken to address the underlaying condition that caused the string to be written to the database, but the presence of the string—the original Triggering condition—remains unaffected.)

Likewise with ACL History Triggers, which react to facts without trying to undo them.

Trigger: Messages

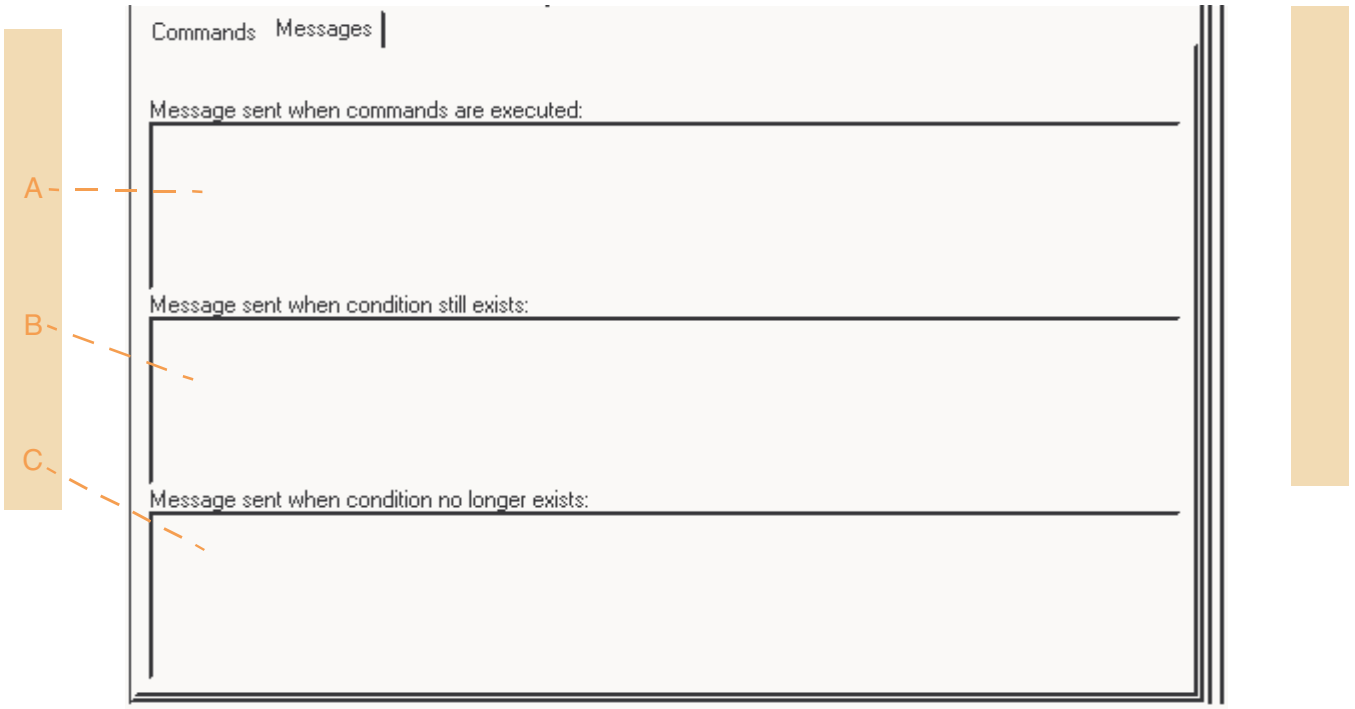

A: text of message sent when Commands are executed

B: text of message sent when the Trigger re-evaluates and the condition still exists

C: text of message sent when the Trigger re-evaluates and the condition no longer exists

Trigger Type: ACL History

Basics ...

Checks for:

- Changes in the ACL History of database(s):
  - any changes
  - changes made by users/servers in a given list
  - changes made by users/servers *not* in a given list

Only those changes that make it into the ACL History Log are checked by Triggers of this type. Any other ACL changes (such as those made on the Advanced tab of the ACL dialog) are not checked, and will not cause ACL History Triggers to fire.

Putting it into practice ...

Users and servers:

This field appears *only* if one of the last two condition options is selected. (Neither of the following fields should be left blank, or the Trigger will not function properly.)

- Option *"Changes made by users/servers in list"*
  - Trigger fires *only* if the detected changes were made by a user/server in the list.
- Option *"Changes made by users/servers not in list"*
  - Trigger fires *only* if the detected changes were made by a user/server *not* in the list.

Example:

Names.nsf plays a vital role in the integrity and operation of any Notes environment. Therefore, it's critical that you be aware of any ACL changes made to this database, since they could result in replication failures.

Monitor for ***any changes***. In addition to e-mailing the Admin, we suggest sending a severity 2 (yellow) IWAlert message to Advanced ServerWatch.

Trigger: ACL History

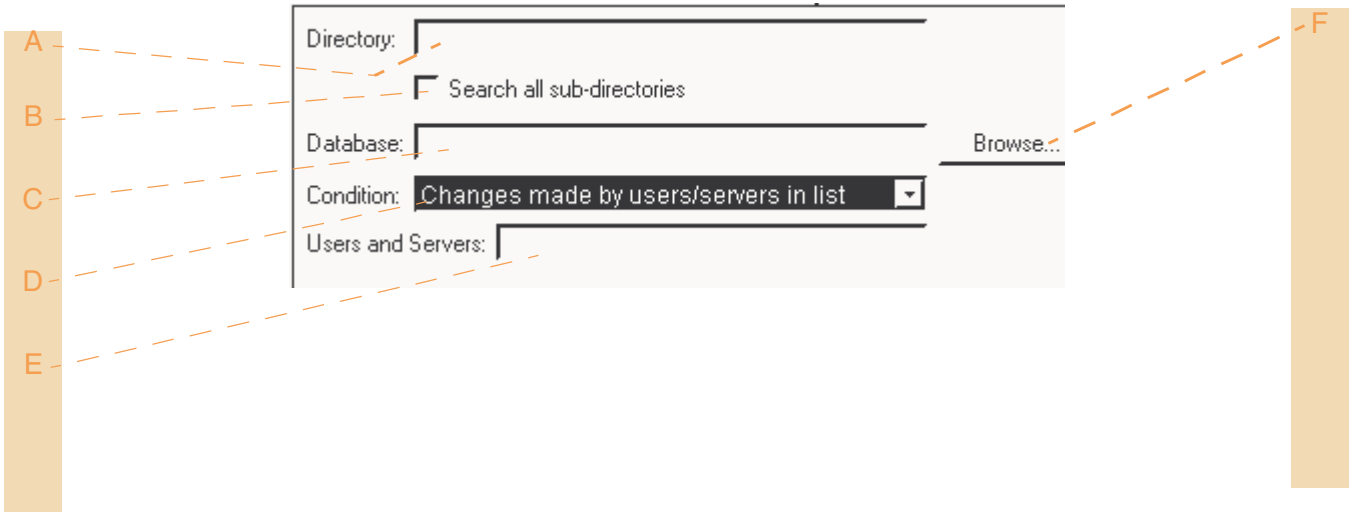

A: location of database(s) to monitor
B: if selected, search all subdirectories
C: database(s) to monitor
D: type of change to check for (double-click icon for options)
E: list of users, servers, and groups (field appears only if changes are to be associated with users in a list, or *not* in a list)
F: button launches Select Database dialog

Trigger Type: Application Trigger

Basics ...

Checks for:

- critical applications at the operating system level
  - Allows admins to restart them, or take other actions as appropriate.
- Domino Server tasks
  - Examples: Router, Replica, Amgr, and so on.
- add-in tasks
  - Examples: gateways, image servers, and so on.
- non-Notes programs running on a Notes server
  - Example: backup utilities.

Putting it into practice ...

Example 1:

You run several instances of Router on Server ABC.

Application Triggers can check not just for the presence of an application, but also for the number of instances currently running. Use this Trigger type to check for the required number of instances of Router.

If the detected number of instances of the Server Task is less than the number specified in the Notes.ini, use IntelliWatch Server Console Commands to launch additional instances.

Example 2:

The Domino HTTP task is particularly subject to problems. IntelliWatch Application Triggers are an effective tool to verify that the task is running, and to restart it if necessary.

Trigger: Application

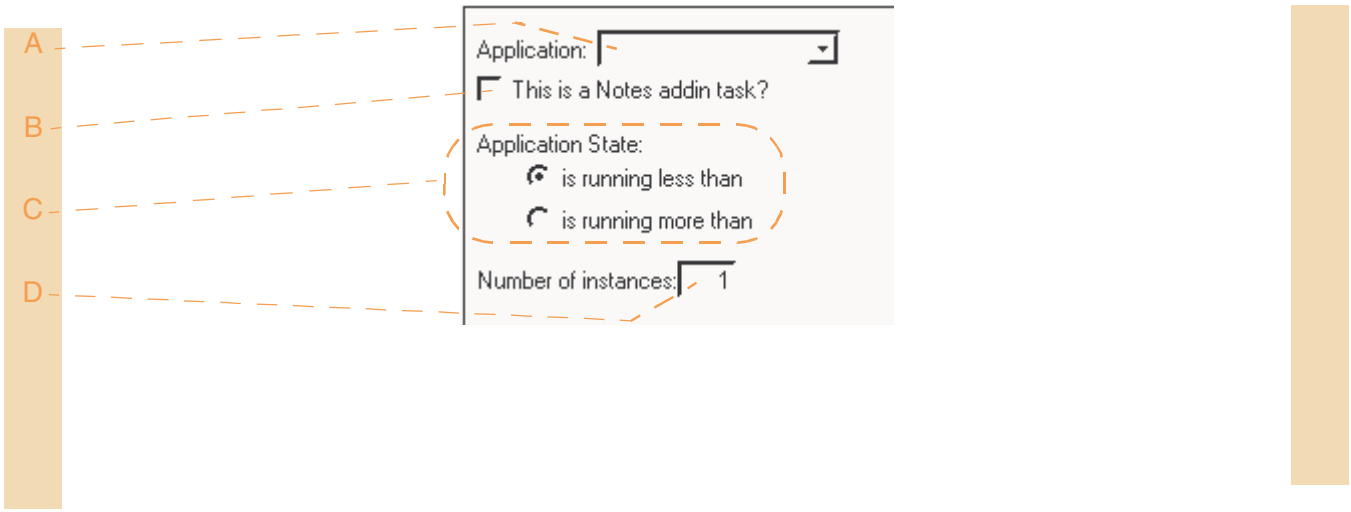

A: name of application to monitor. Click on down arrow of combo box to access list, or type in by hand (double-click icon for examples)

B: application being monitored is Notes server add-in task (if selected)

C: number of instances to check for is < or > number listed in Number of Instances field (see D)

D: number of instances of monitored application that should be running

Trigger Type: Availability Trigger (Service)

Basics ...

Checks for:

- NT Services
- response threshold for TCPIP applications
  - Examples: HTTP, LDAP, IMAP, SMTP. Checks by pinging the port.
- SMTP MTA Mail (using Mail Probe)

Mail Probe monitors:

a) SMTP MTA, by verifying that mail is being routed to the internet.

b) How long mail takes to get from a mail server to the server running the SMTP MTA.

c) How long it takes to get a response forwarded back from an internet address.

Requirements:

- Internet address that can echo (or forward) a response back to the sender
- dummy mail user in the NAB with a mail file on the server

Putting it into practice ...

Example:

You use SMTP MTA to route mail to the Internet, and want to keep a close eye on its performance. The Service Option of Availability Triggers can help you do this, by means of a Mail Probe.

The Mail Probe sends an E-mail to the address given, and the specified mail file receives trace reports.

Caveats:

The Mail Probe does not work in all circumstances. The success of the probe—and/or the validity of the results returned—can be affected by your ISP, the address used to forward mail, and other factors.

Since these factors are beyond the control of both the IntelliWatch program and IBM support, the ultimate success of the Mail Probe cannot be guaranteed.

In situations where your ISP does not meet the Mail Probe's requirements, IBM support may be able to provide a working alternative:

The TCPIP option checks for port availability, NOT to see if the port is actually being used.

For example, the default port for SMTP is 25, and entering either value in the Port field will cause the Trigger to check for that port's availability. But, as long as the port is available, the Trigger will not fire, even if the SMTP task is not loaded.

Trigger: Availability Trigger (Service)

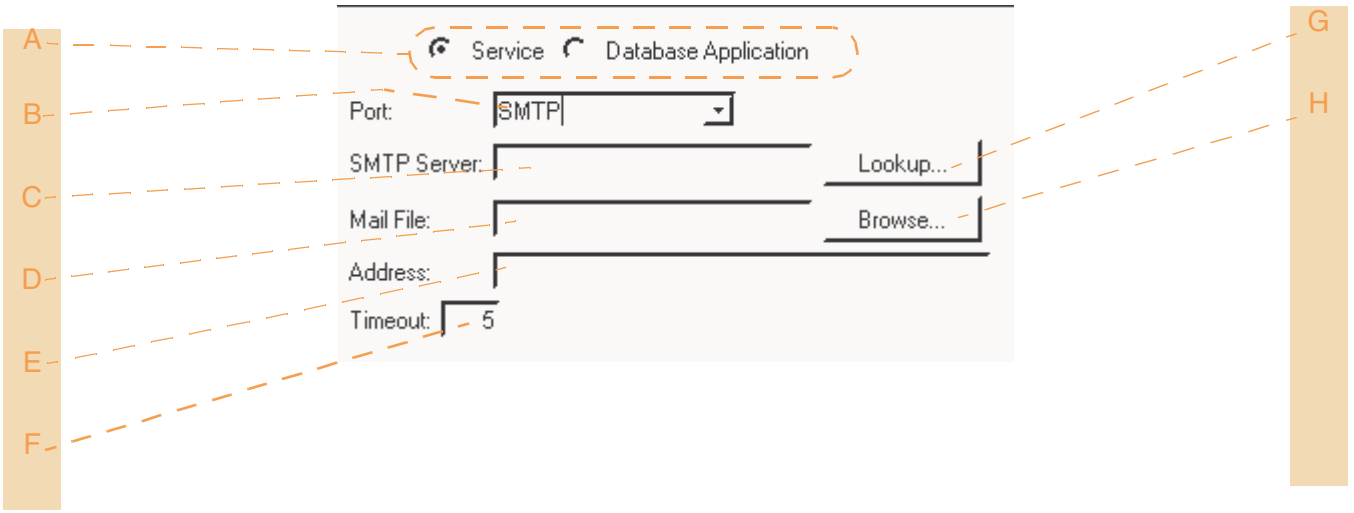

A: item to monitor

B: port to monitor. Click on down arrow of combo box to access list, or type in port by hand (double-click icon for examples)

C: server running SMTP MTA (mail probe). Select using Lookup button, or type in by hand.

D: name of valid mail file (for mail probe). Select using Browse button, or type in by hand.

E: internet address which can echo/forward response to sender

F: wait (in seconds) before resource deemed unavailable

G: button launches Select Server dialog

H: button launches Select Database (mail file) dialog

Trigger Type: Availability Trigger (Database Application)

Basics ...

Checks for:

- database Applications
  - Example: response time of Notes databases.

Allows Admins to assign database applications a custom time-out (at which point the resource will be deemed unavailable, and configured actions will be taken.

Putting it into practice ...

Example:

Acme Accounts Payable uses **CriticalTransactions.nsf** for Transaction Processing. When access to this database is too slow, Acme is unable to fulfill their SLAs with customers, leading to loss of business.

Acme's IT team configured an Availability Trigger to fire whenever a given (access) time-out for this database was exceeded (see field E on the following page).

When the threshold was exceeded, the Trigger fires, and all selected Commands (such as notification) were executed.

To be in a better position to analyze the application's performance, they also had Monitor generate a Trigger statistic (call *CriticalTrans.ThresholdExceeded*), and used another (Statistic) Trigger to monitor its value, and notify them when the value became critically high.

Trigger: Availability Trigger (Database Application)

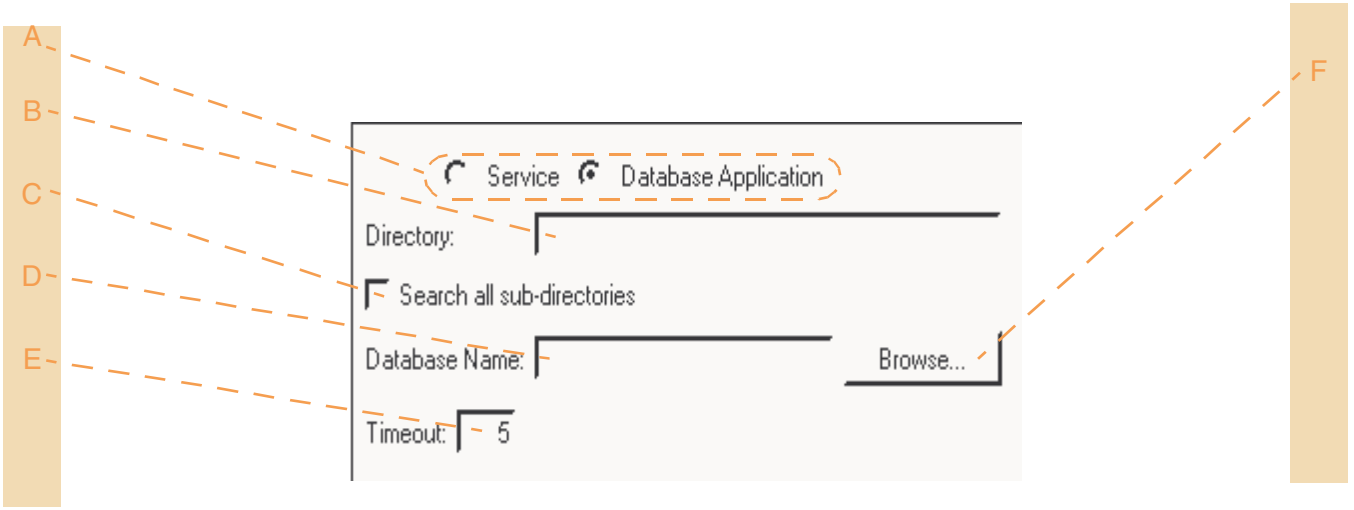

A: item to monitor
B: full path to directory containing database(s). Leave blank, or enter [Directory] for the Notes data directory
C: if selected, search all subdirectories
D: comma-delimited list of databases. Leave blank (or use *.*) for all. Select databases using the dialog activated by the button at letter F, or type them in by hand.
E: enter an integer to specify a timeout (in seconds)
F: button launches Select Database dialog

Trigger Type: Compound Trigger

Basics ...

Checks for:

- two conditions at the same time
  - Conditions may be of the same or different types.

As you plan and configure Compound Triggers, please remember that *both* specified conditions must be true for the Trigger to fire.

Environment variables are handled differently by Compound Triggers: a number is appended to the variable name, to specify which part of the Compound Trigger uses the variable.

For instance, if the Compound Trigger contains an Application Trigger (part 1) and a File Trigger (part 2) the variables set by the Trigger would be <APPLICATION1> and <FILE2>.

Putting it into practice ...

Condition-specific dialogs:

The initial dialog (see following page) is the same for all Compound Triggers. Clicking on one of the two Edit Condition buttons, however, launches a different, type-specific configuration dialog, based on the selection made in the associated combo box.

In all, there are thirteen different dialogs, each with a unique combination of parameters.

Since the information required by most dialog fields is self-explanatory, usage tips for each type are not included here. Instead, illustrations of each dialog are included, with explanations on the same page for fields requiring additional clarification. For information on a given condition type, consult the section for the relevant (Trigger) type.)

Example:

The HTTP task seldom fails at Island ISP, Inc., but it's response time is often unacceptably slow.

Admins set up a Compound Trigger combining an Application condition (that checked to see that the task was running), and Availability condition (that monitored its response time). If response time exceeded 30 seconds, notification actions were taken.

Trigger: Compound Trigger

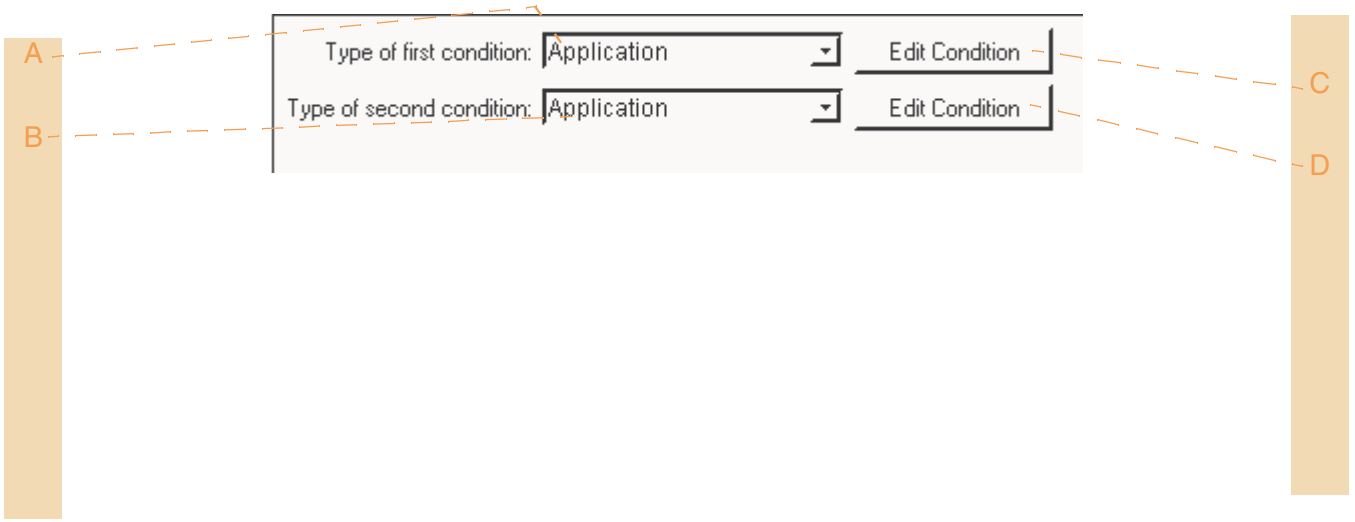

A: type of first condition or Trigger to be monitored. Click on down arrow of combo box to access list, or type in port by hand (double-click icon for examples).

B: type of second condition or Trigger specified to monitor

C: clicking this button allows configuration of the first condition

D: clicking this button allows configuration of the second condition

Trigger Type: Compound Trigger: Edit Condition dialog

Basics ...

Either Edit Condition button launches the same dialog for a given condition type. Most dialog fields are self-explanatory. Details are provided where necessary.

For field definitions of the various condition types, see *Compound Trigger Condition: on page 85* through *Compound Trigger Condition: on page 97*.

Putting it into practice ...

Compound Trigger Condition: ACL History

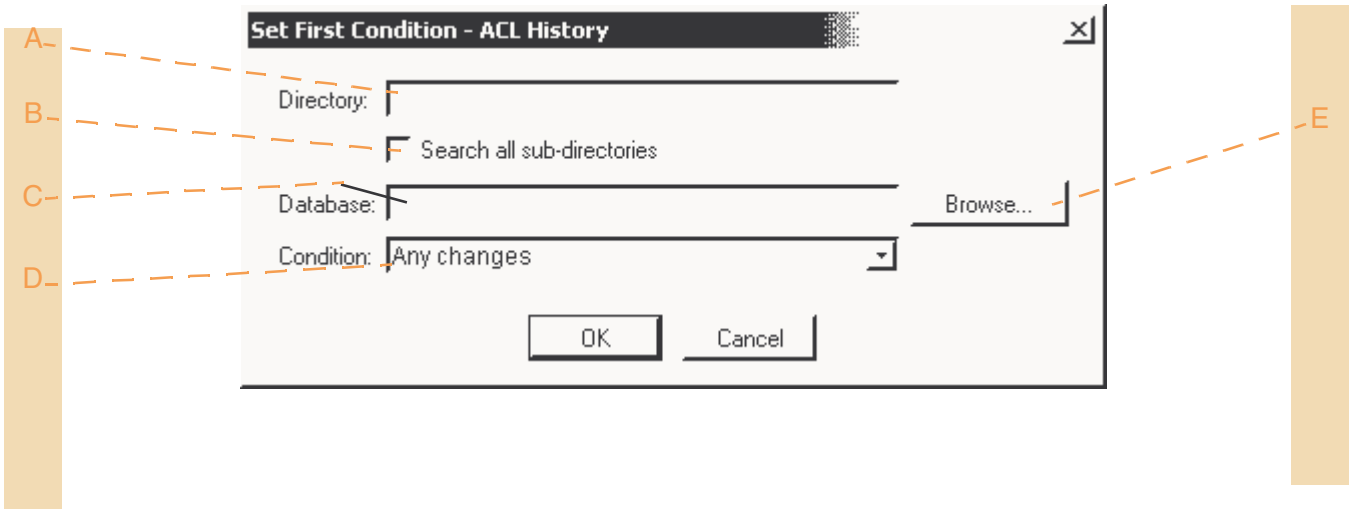

A: full path to directory containing database(s). Leave blank, or enter [Directory] for the Notes data directory.
B: if selected, search all subdirectories
C: comma-delimited list of databases. Leave blank (or use *.*) for all.
D: condition for which to monitor. Click on combo box arrow to access list (double-click icon for options).
E: button launches Database Selection dialog

Compound Trigger Condition: Application

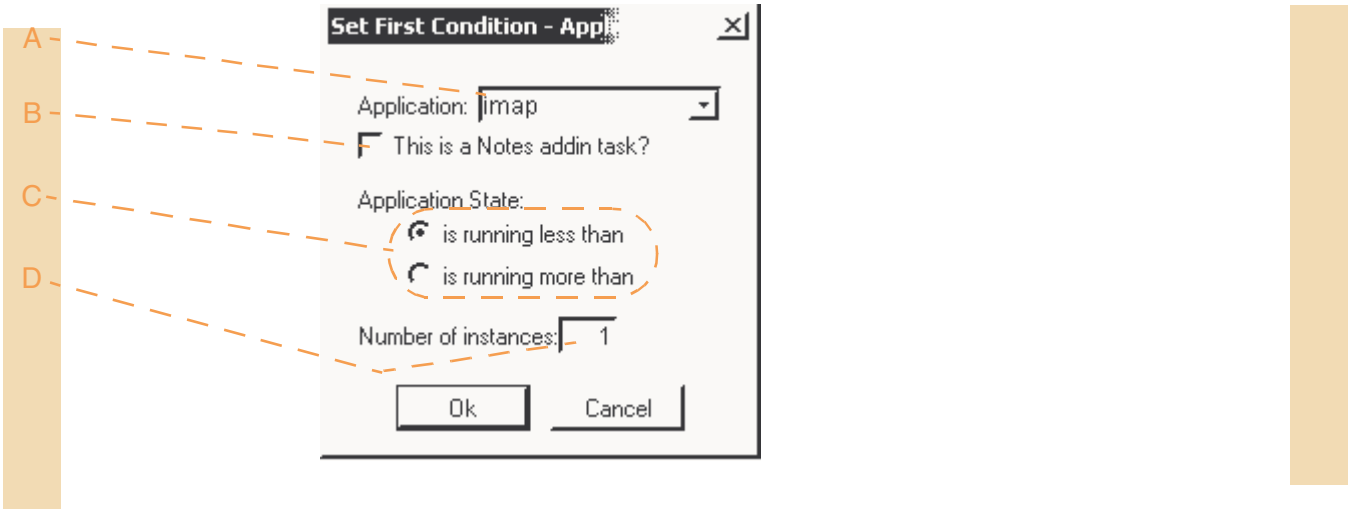

A: type of Application to be monitored. Click on down arrow of combo box to access list, or type in application by hand (double-click icon for options).

B: if selected, application is Notes add-in task

C: specify < or > operator (used in conjunction with field at D to determine if condition has been met

D: number of instances of monitored application that should be running

Compound Trigger Condition: Availability (Database Application)

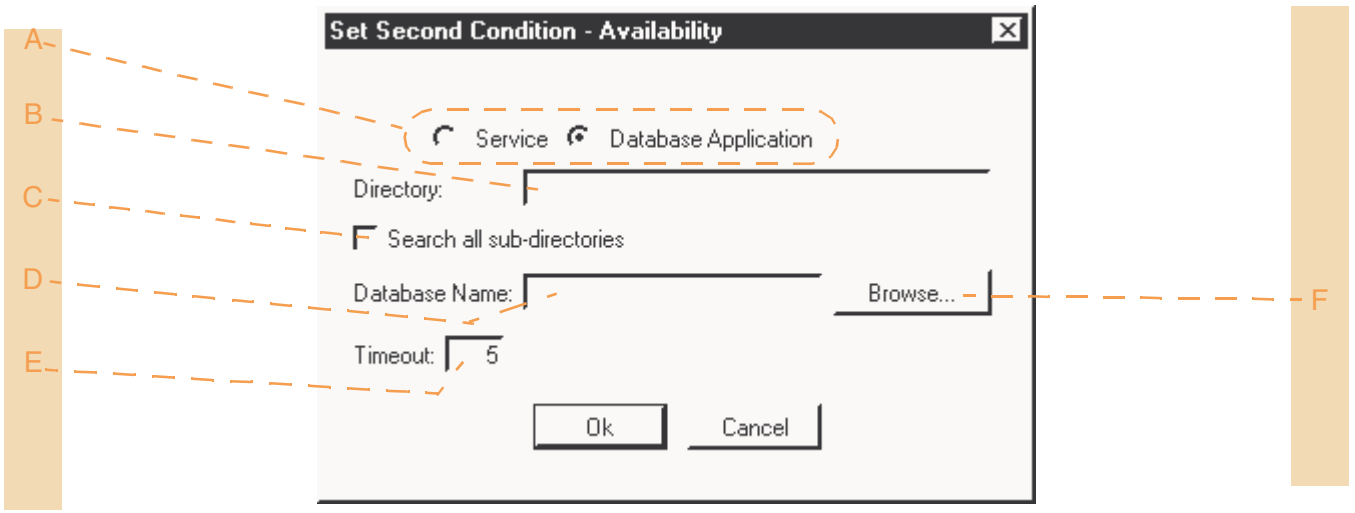

A: item to monitor
B: full path to directory containing database(s). Leave blank (or enter [Directory] for the Notes data directory.
C: if selected, search all subdirectories
D: comma-delimited list of databases. Leave blank (or use *.*) for all. Select databases using the dialog activated by the button at letter F, or type it in by hand.
E: enter an integer to specify a timeout (in seconds)
F: button launches Select Database dialog

Compound Trigger Condition: Availability (Service)

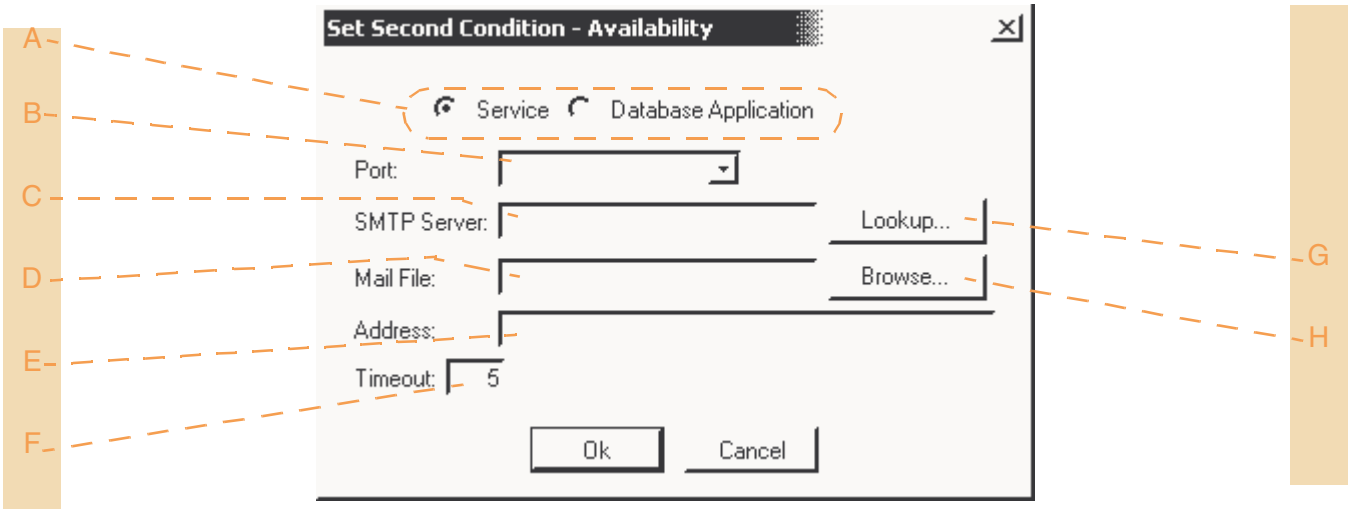

A: item to monitor
B: port to monitor. Click on down arrow of combo box to access list, or type in port by hand (double-click icon for examples).
C: server running SMTP MTA (mail probe). Select using Lookup button, or type in by hand.
D: name of valid mail file (for mail probe). Select using Browse button, or type in by hand.
E: internet address which can echo/forward response to sender
F: wait (in seconds) before resource deemed unavailable
G: button launches Select Server dialog
H: button launches Select Database dialog

Compound Trigger Condition: Database Activity

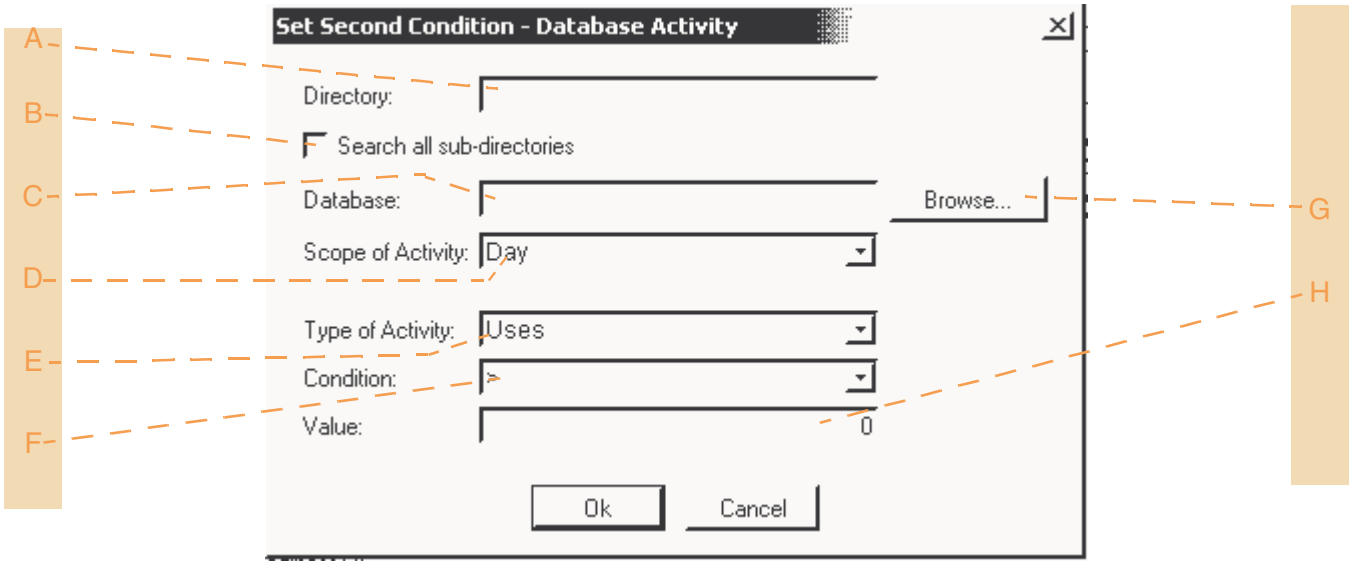

- A: full path to directory containing database(s). Leave blank (or enter [Directory] for the Notes data directory.
- B: if selected, search all subdirectories
- C: comma-delimited list of databases. Leave blank (or use *.*) for all. Select databases using the dialog activated by the button at letter G.
- D: scope of activity: day, week, month
- E: type of activity for which to monitor: uses, reads, writes
- F: select condition type. Click on combo box arrow to access list (double-click icon for options).
- G: button launches Database Selection dialog
- H: value required to cause Trigger to fire

Compound Trigger Condition: Database Scan

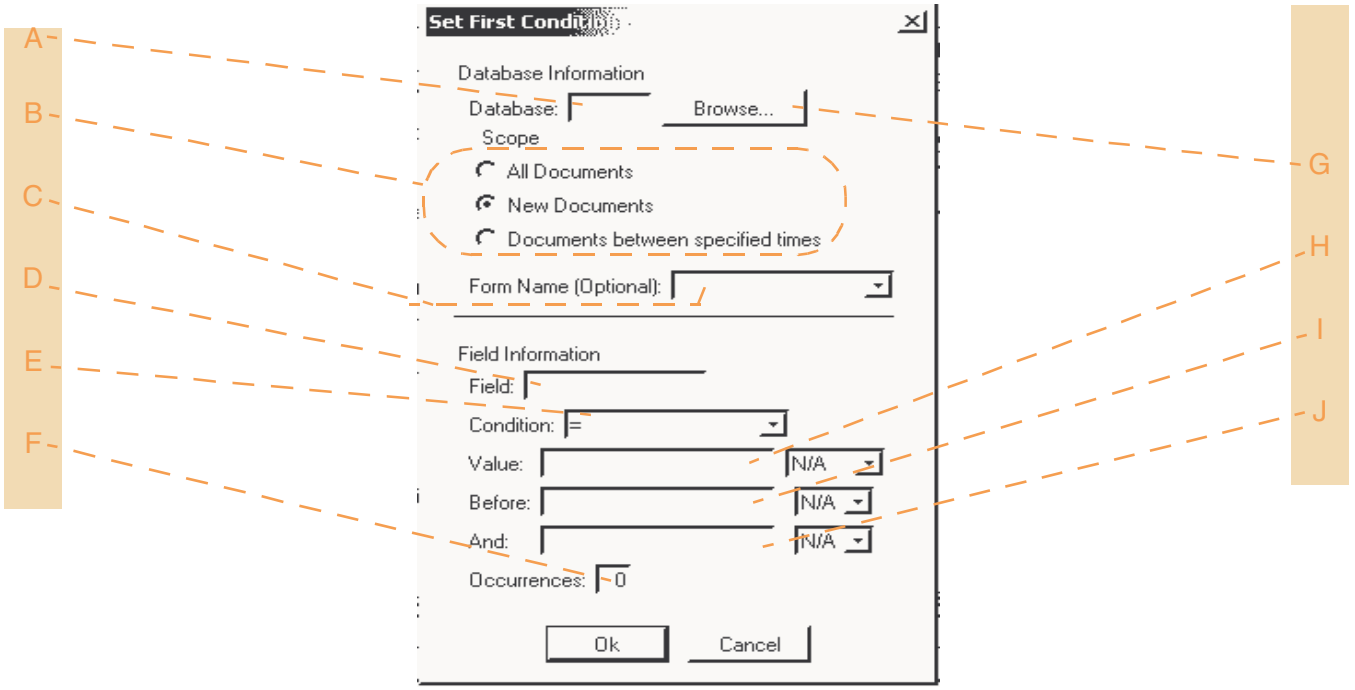

A: comma-delimited list of databases. Leave blank (or use *.*) for all. Select databases using the dialog activated by the button at letter G.

B: select desired option to specify scope of scan

C: optional name of form to be monitored. Click on down arrow of combo box to access list, or type in application by hand (double-click icon for examples)

D: field to monitor

E: condition operator (double-click icon for options)

F: number of occurrences required to cause Trigger to fire

G: button launches Select Database dialog

H: value to search for (double-click icon for units)

I: lower limit (when condition operator is Range). End point values are included. (See H for units.)

J: upper limit (when condition operator is Range). End point values are included. (See H for units.)

Compound Trigger Condition: Document Count

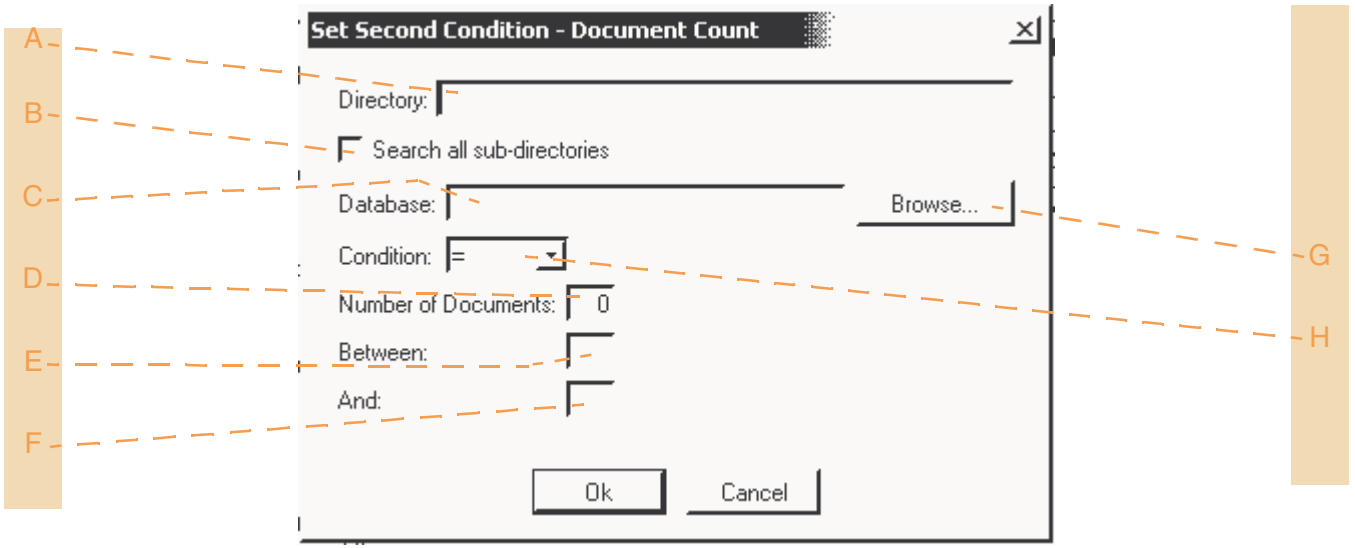

A: full path to directory containing database(s). Leave blank, or enter [Directory] for the Notes data directory.
B: if selected, search all subdirectories
C: comma-delimited list of databases. Leave blank (or use *.*) for all. Select databases using the dialog activated by the button at letter G.
D: specific number of documents to check for
E: lower value of range (end point values are included in the range). Field applies only when Range operator is selected.
F: upper value of range (end point values are included in the range). Field applies only when Range operator is selected.
G: button launches Database Selection dialog
H: operator governing Condition field D (double-click icon for options)

Compound Trigger Condition: Document Time-out

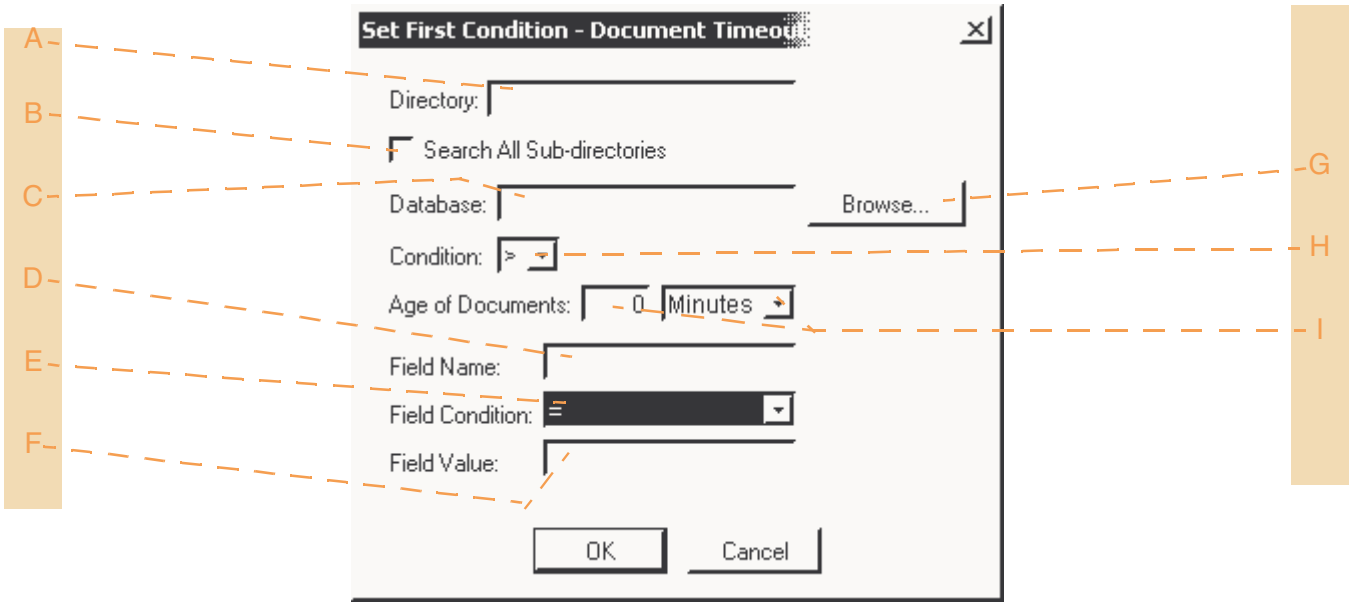

A: full path to directory containing database(s). Leave blank, or enter [Directory] for the Notes data directory.
B: if selected, search all subdirectories
C: comma-delimited list of databases. Leave blank (or use *.*) for all. Select databases using the dialog activated by the button at letter G.
D: name of field which must meet following condition
E: operator governing field condition. Click on combo box arrow to access list (double-click icon for options).
F: value required for condition to be met
G: button launches Database Selection dialog
H: condition specifier < or >
I: enter numerical value in text field, and select units of time using combo box (minutes, hours, or days)

Compound Trigger Condition: Replication Integrity

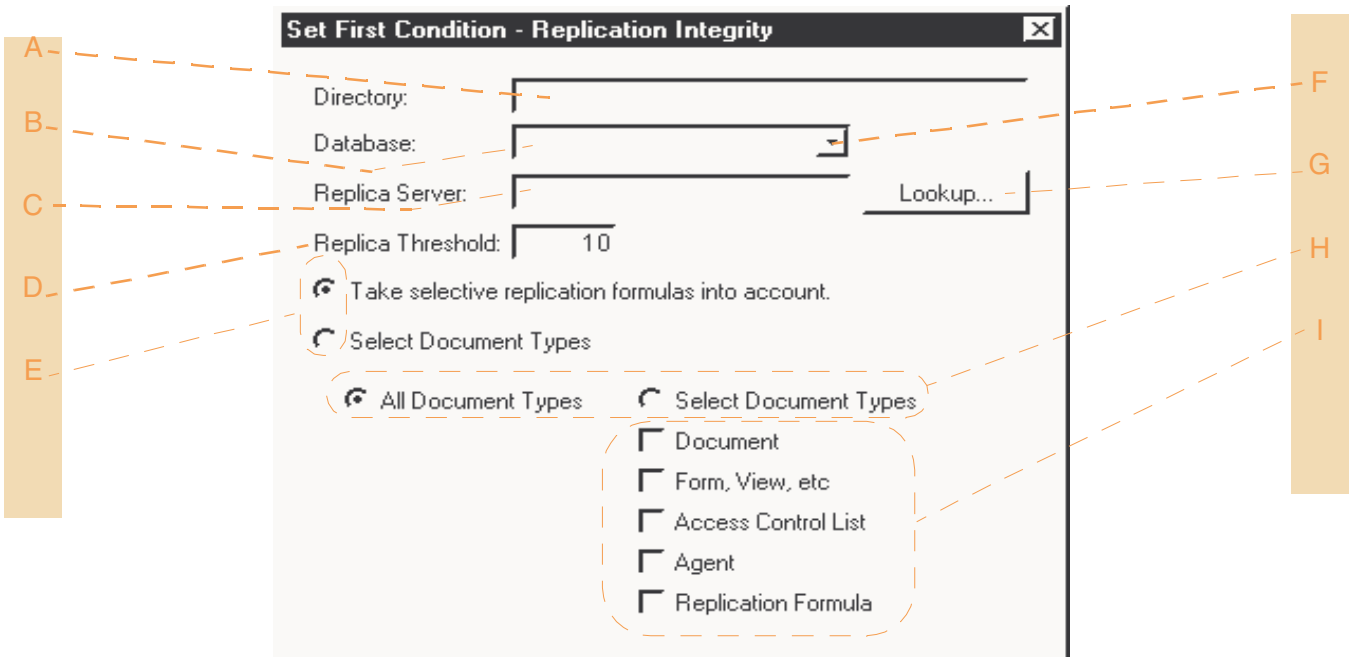

- A: full path to directory containing database(s). Leave blank (or enter [Directory] for the Notes data directory.
- B: database to examine
- C: select Replica Server
- D: set threshold for minimum number of differences required for Trigger to fire
- E: elect EITHER to use replication formulas, or to limit the check by document types
- F: select database using combo box, or type in name (field uses auto-complete for known databases)
- G: button launches Server Selection dialog
- H: choose EITHER all types, or use check boxes to limit Trigger to certain types
- I: document types available as selection criteria (multiple selection allowed)

Compound Trigger Condition: Replication Readiness

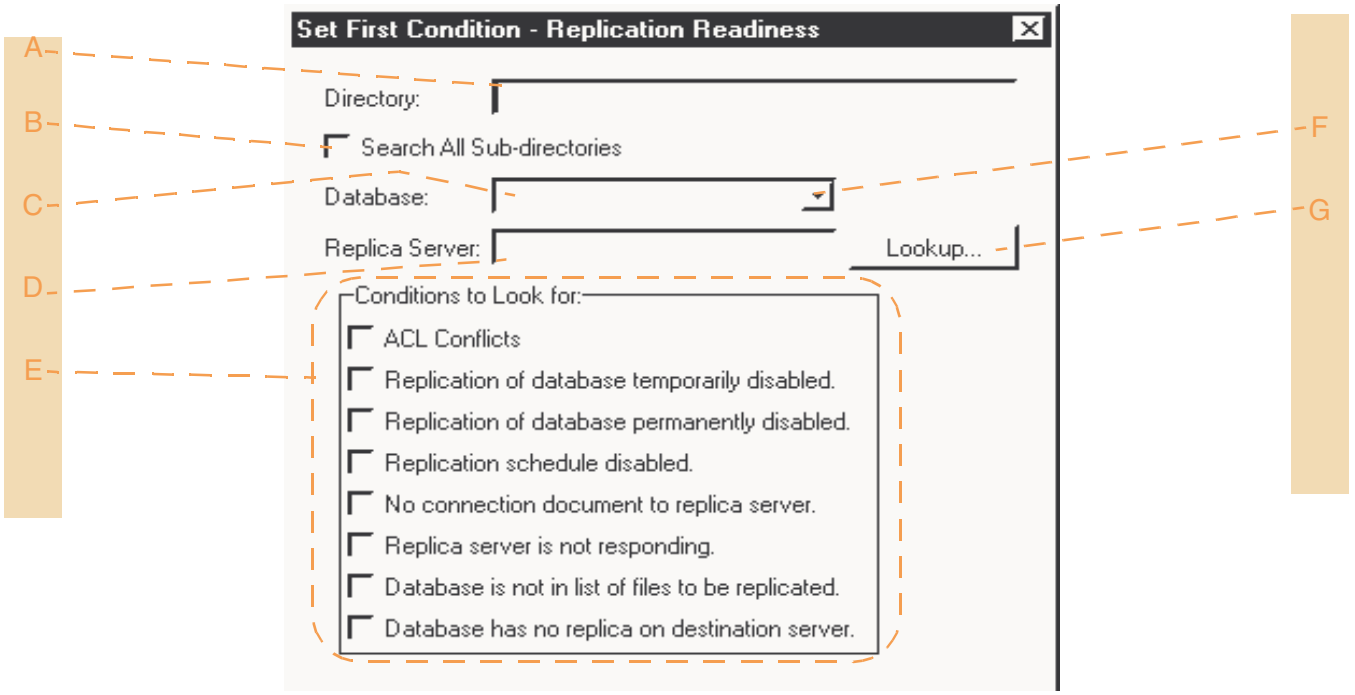

A: full path to directory containing database(s). Leave blank, or enter [Directory] for the Notes data directory.
B: check to search sub-directories
C: select database to examine
D: select Replica Server
E: select all desired conditions affecting replication
F: use combo box to select database
G: button launches Server Selection dialog

Compound Trigger Condition: Statistic

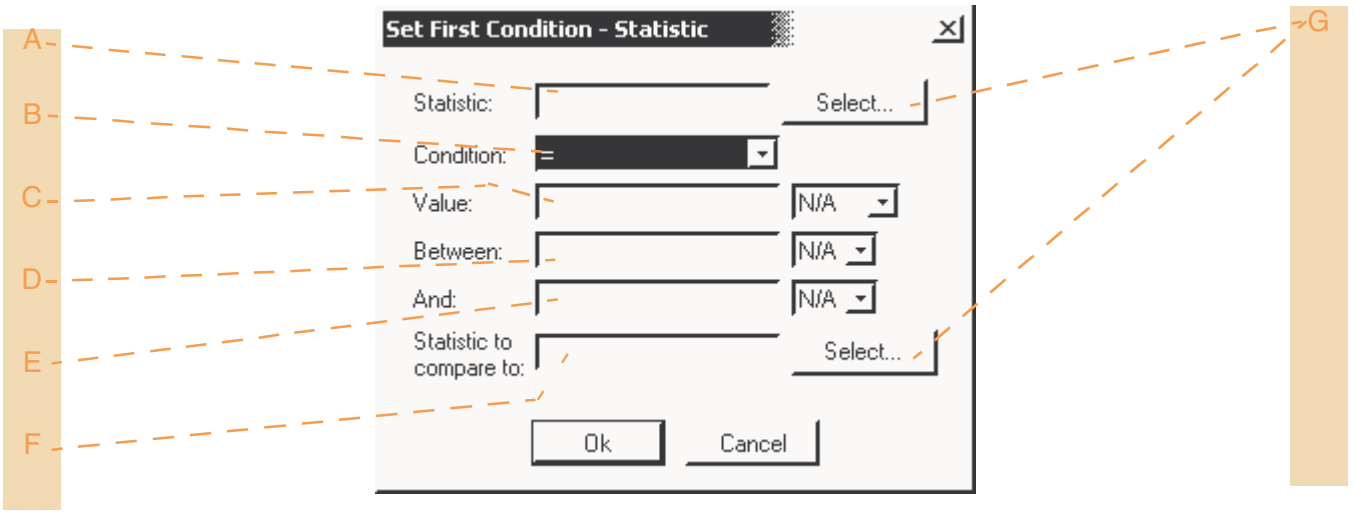

A: select Statistic to be monitored by clicking on button at G

B: select condition. Click on arrow of combo box to access list (double-click icon for options).

C: value required for Trigger to fire. Click on the arrow of the combo box to access the list of units (double-click icon for units).

D: lower limit (when condition operator is Range). End point values are included. (See C for units.)

E: upper limit (when condition operator is Range). End point values are included. (See C for units.)

F: second Statistic to be checked as a percentage of the first Statistic (in field A). This field is enabled *only* when % is selected as the unit of Value.

G: button launches Select Statistic dialog

Compound Trigger Condition: User

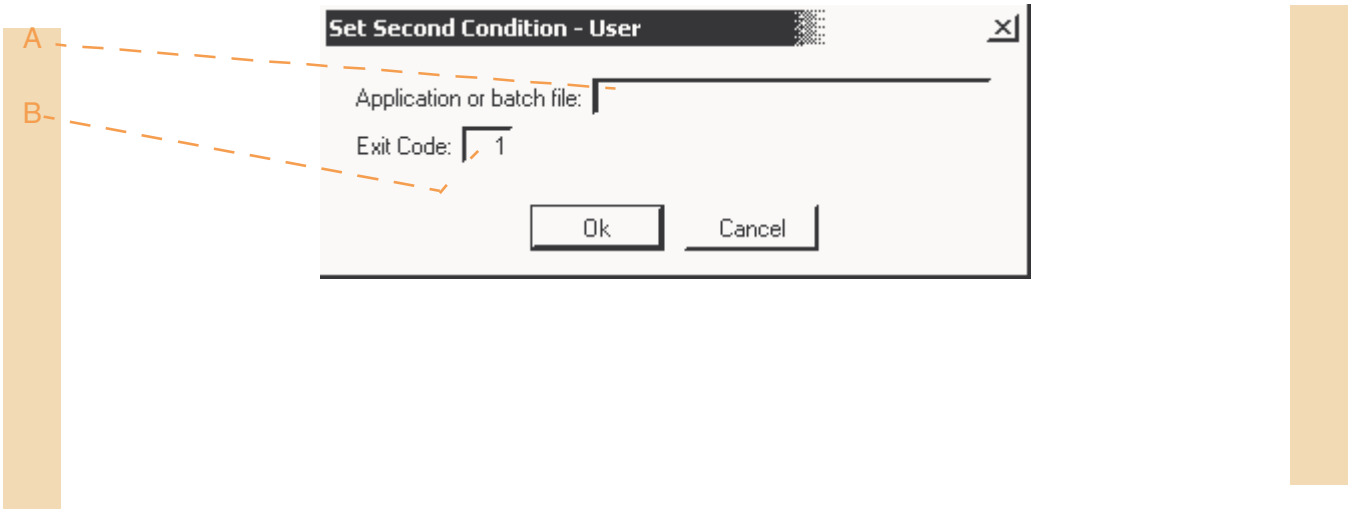

A: name of the application or batch file to be run, followed by any command-line arguments
B: exit code which meets condition

Compound Trigger Condition: White Space

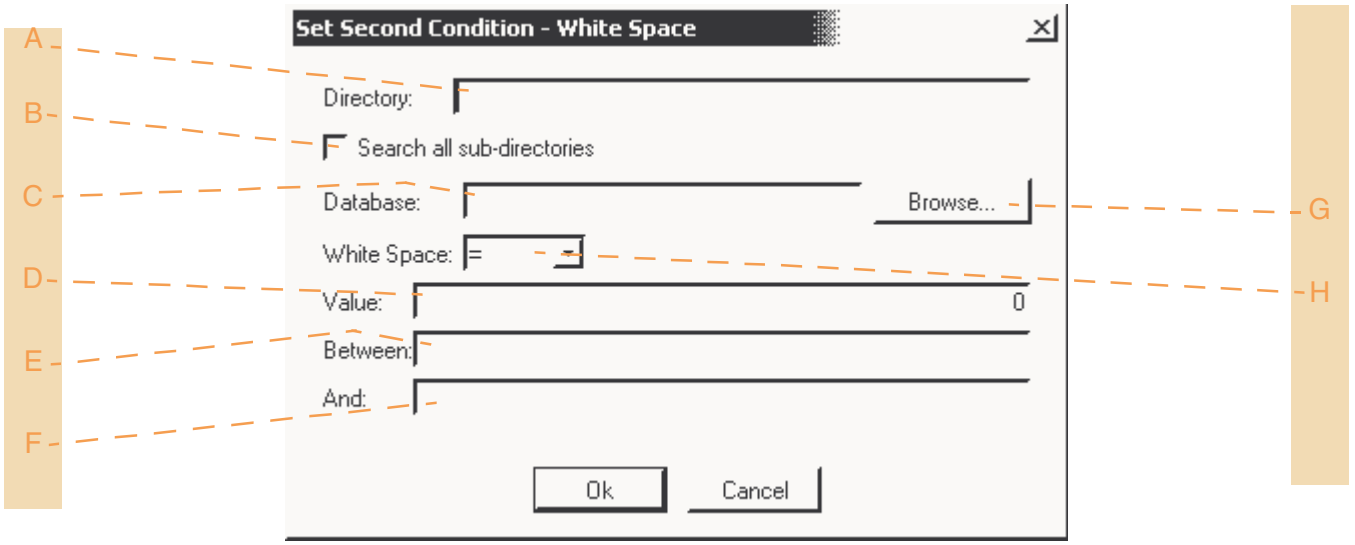

A: full path to directory containing database(s). Leave blank (or enter [Directory] for the Notes data directory.
B: if selected, search all subdirectories
C: comma-delimited list of databases. Leave blank (or use *.*) for all. Select databases using the dialog activated by the button at letter G.
D: specific value (%) of white space to check for
E: lower limit (when condition operator is Range). End point values are included.
F: upper limit (when condition operator is Range). End point values are included.
G: button launches Database Selection dialog
H: operator governing white space value field D (double-click icon for options)

Trigger Type: Database Scan Trigger

Basics ...

Checks for:

- user-specified text string
  - Checks specified field in any Notes database.

Search for occurrences in:

- all documents
- new documents
- documents between specific times

New in Maintenance Release 27.33 is the keyword <DATABASE_OCCURRENCES_FOUND>, which returns the number of times the search string was actually found (as opposed to the number stated on the Condition tab of the Trigger).

Putting it into practice ...

Form:
Use to search for a value in multiple views.

Field:
Required parameter.

Condition:
Contains and **does not contain** for textual fields.

=, <>, >, < and **range** for numeric fields.

Value:
Make as specific as you can, to avoid "false positives".

Between/And:
Useful for troubleshooting.

Occurrences:
Reduce the number of notifications you receive by increasing the Occurrences field value. For example, if Occurrences equals 3 for the 'Unable to Replicate With Server' message, notification occurs only if three error messages are found *in one monitoring cycle*.

Example:

Database Triggers can be used to key off texts that *aren't* found, as well as those that are. Consider this real-world example:

The index to one of White Incorporated's critical databases occasionally became corrupt, and this corruption was propagated across the Domino network when scheduled replication took place, leading to system-wide down time while the issue was addressed.

Whenever this issue arose, the text string "possible index corruption in KeyDb.nsf" appeared in the Notes log.

A Database Trigger was configured to periodically search for this string, and to initiate Replication when it was *not* encountered.

A second Database Trigger was configured to notify Admins when the string *was* encountered.

Trigger: Database Scan Trigger

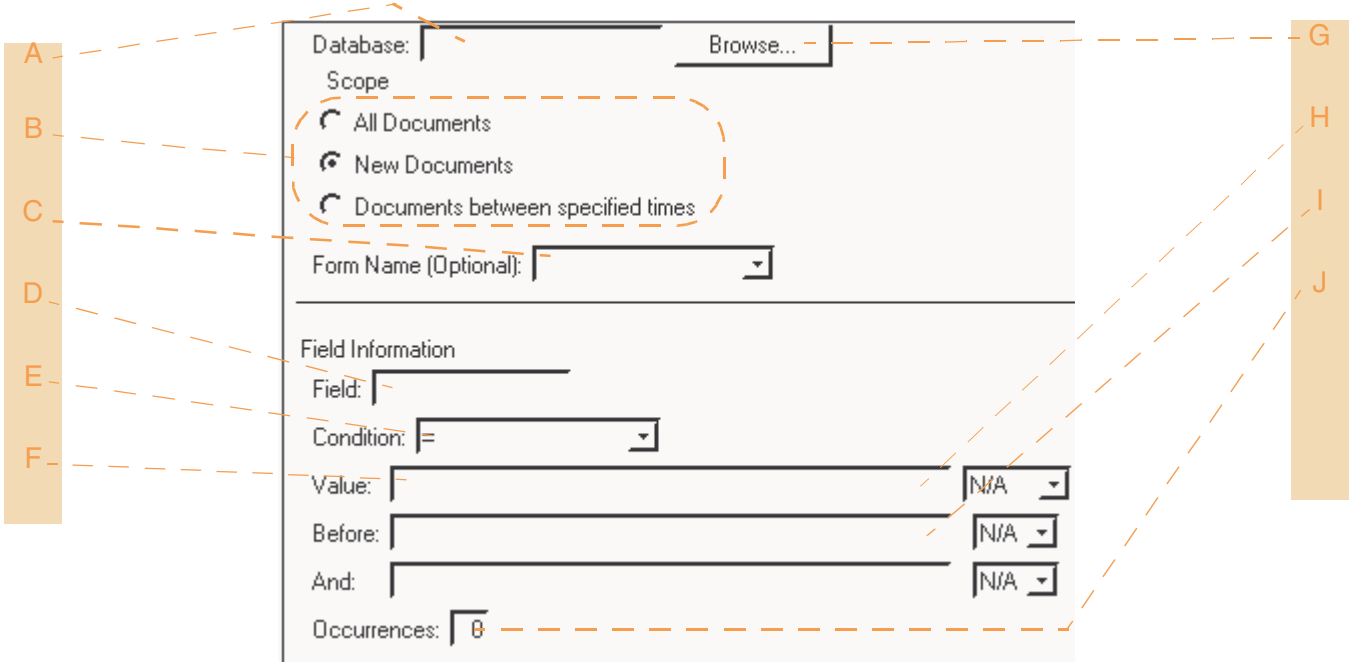

A: name of database to monitor
B: subset of documents to monitor
C: optional: limits search to given form name (double-click icon for examples)
D: name of field to monitor
E: operator type for condition statement (double-click icon for options)
F: value to search for in specified field (double-click icon for options)
G: button launches Select Database dialog
H: lower limit (when condition operator is Range). End point values are included. (see E for units.)
I: upper limit (when condition operator is Range). End point values are included. (see E for units.)
J: number of occurrences of database condition required to trigger actions

Trigger Type: Database Activity

Basics ...

Checks for:

- Uses
- Reads
- Writes

Use to identify databases that are used infrequently—if at all. (You may wish to remove these databases from your system, to conserve resources.)

This Trigger type can check for activity for the previous day, week, or month.

When specifying Uses, Reads or Writes in a Trigger, bear in mind that Uses do not require user activity (think of database compaction, for example). Consult your Notes documentation for details of what constitutes, Uses, Reads and Writes, then configure the Database Activity Trigger in accordance with the kind of activity you're trying to identify.

Putting it into practice ...

Example:

John Smith left Widget Enterprises last month, but the Admin responsible for removing the mail files of former employees forgot to perform the requisite actions.

A Database Activity Trigger identified this database as having had no Reads during the previous month. Admins were alerted, and the database removed from the system.

Trigger: Database Activity

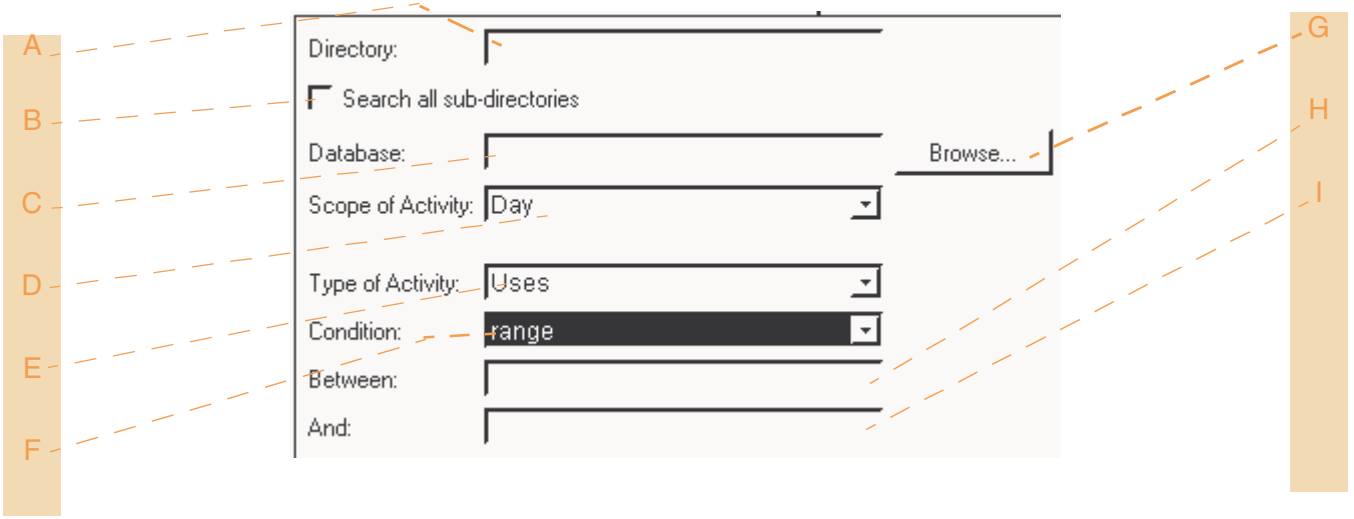

A: location of database to monitor
B: if selected, search all subdirectories
C: database to monitor
D: scope of activity to monitor (day, week or month)
E: type of activity to monitor (uses, reads or writes)
F: operator type for condition statement (double-click icon for options)
G: click Browse to select the database from a list
H: lower limit (when condition operator is Range). End point values are included.
I: upper limit (when condition operator is Range). End point values are included.

Trigger Type: Document Count

Basics ...

Checks for:

- empty databases
- number of documents in specified Notes database

Care should be taken to limit the monitoring frequency of Document Count Triggers. Depending on the terms of any applicable SLAs, you may want to configure such a Trigger for daily (or weekly) monitoring, rather than, say, hourly. The presence of many large databases can cause execution of this Trigger type to take a very long time.

Putting it into practice ...

Example:

Databases that contain few (or no) documents can clutter up your environment. To identify all empty databases on servers running Monitor, create a Document Count Trigger, using the following parameters (and replicate it to all servers where you want to perform the check):

- check "Search all sub-directories"
- database: *.nsf
- condition selected: =
- number of documents: 0

Use the <DATABASE> keyword in messages sent by Document Count Triggers to notify you which databases are empty (or contain less than X-number of documents).

With the <DATABASE> keyword in the message field of the Trigger, each time a database is encountered that fits the configured search profile, the name of that database is passed to any notification commands sent by the Trigger.

Once identified, you can remove these empty databases to reduce unnecessary replication—especially over WAN.

Trigger: Document Count

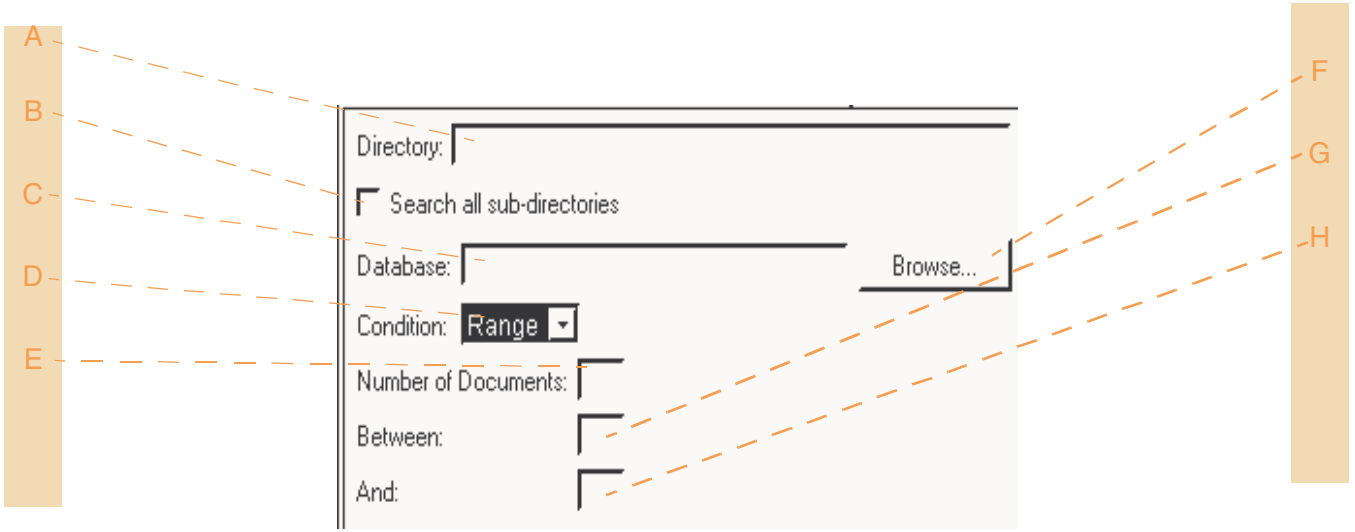

- A: location of database whose documents are to be counted
- B: if selected, search all subdirectories
- C: database(s) on which to execute document count
- D: operator type for condition statement (double-click icon for options)
- E: number of documents for which to monitor (not active here because of Range operator)
- F: button launches Select Database dialog
- G: lower limit (when condition operator is Range). End point values are included.
- H: upper limit (when condition operator is Range). End point values are included.

Trigger Type: Document Time-out

Basics ...

Checks for:

- age of documents in a database
  - Can be coupled with a field condition (see below).

Field condition options:

Assumes that the Time-out condition has been met.

- =
  - Trigger fires *only* if field value is an *exact* match.
- <>
  - Applies to numerical fields.
- Exists (Field Value does not apply)
  - If the field is found, the Trigger will fire.
- Does not exist (Field Value does not apply)
  - If the field is *not* found, the Trigger will fire.
- Contains
  - Trigger fires if the target string is found.
- Does not contain
  - Trigger fires if: 1) the target string is found; 2) the named field does not exist.

Putting it into practice ...

Computed vs editable fields:

When selecting a field to search for a value, ascertain if the field is *computed* or *editable*, then consult your Lotus documentation as to the differences in functionality between these field types.

Condition options:

For all condition options, the field specified must be part of the document's Summary Information.

For the four text-comparison options, rich-text fields cannot be used.

Example:

At Deluxe Lighting, Inc., certain mail users have a tendency to allow old documents to accumulate in their mail files, taking up valuable disk space.

Admins at Deluxe configured a Document Time-out Trigger to notify them (and the USER) whenever such databases are encountered. They configured the message sent at notification so it includes suggestions as to the appropriate steps for users to take to remedy the situation.

Trigger: Document Time-out

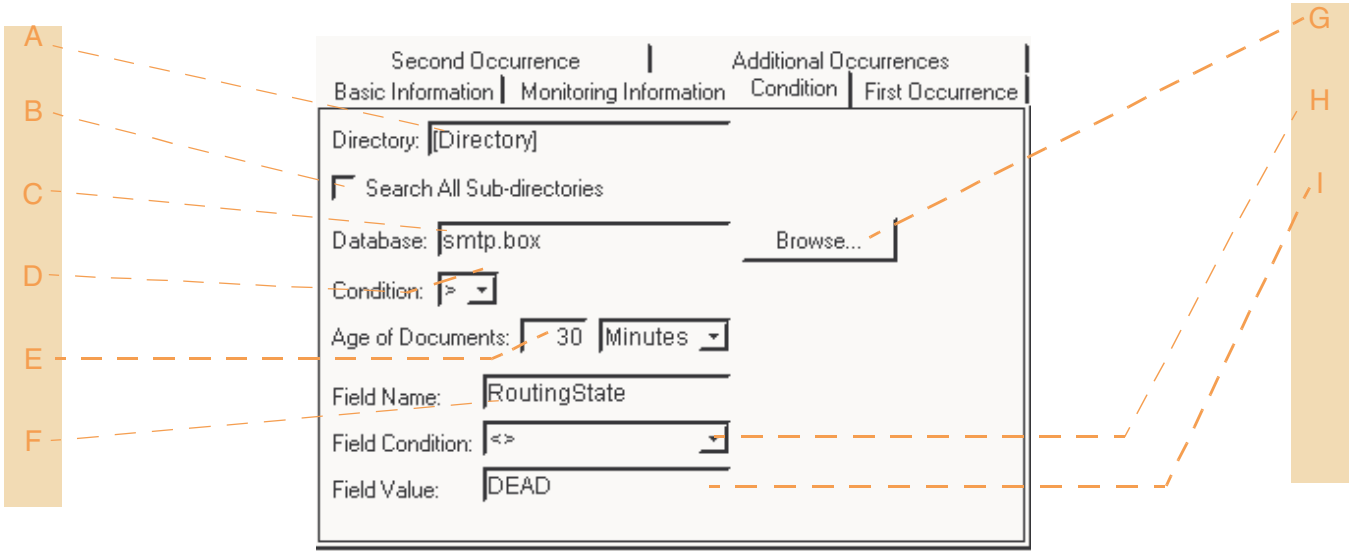

- A: location of database to monitor
- B: if selected, search all subdirectories
- C: database to monitor (supports wildcard entries such as *.nsf)
- D: operator type for condition statement (< or >)
- E: age of documents in units specified (minutes, hours, or days)
- F: field name which must meet Field condition (see H)
- G: button launches Select Database dialog
- H: operator type for field condition (double-click icon for options)
- I: value that field should or should not equal

Trigger Type: File Trigger

Basics ...

Checks for:

- file information
  - notifies of user-specified conditions
- when database exceeds user-specified size
- changes made to configuration files
  - Examples: notes.ini and desktop.dsk files.

Use wildcards to monitor multiple files. For example, monitor all mail files by specifying /notes/mail/*.nsf.

=, <>, >, < and **Range** operators apply to numeric fields.

If searching for Date/Time value, use the syntax:

M/d/yy h:mm

Condition type field:

Depending on the condition type selected, certain fields are enabled/disabled. If Date/Time is selected, the combo boxes governing units are disabled (the condition doesn't require them).

Condition field:

Similarly, the "Between"/"And" fields are only enabled if the Condition is Range.

Putting it into practice ...

On clustered servers:

On a clustered server, if a File Trigger is set up to notify users if their db is too large, dead mail may result, since IntelliWatch won't find users for databases that have another server as their home server (even though the database still exists on this server in case of failure).

To avoid the dead mail: Create a mail user called "NOT FOUND" - all mail will be sent to this user if IW cannot resolve the name (the env var "MailUser" is set to "NOT FOUND" if the lookup fails. Set "Remove documents not modified in last X days...." low to keep the db small.

Update Date/Time or bytes field after triggering:

Field must be a reference to a configuration parameter:

– in the registry (NT)

```
HKEY_LOCAL_MACHINE/SOFTWARE/Candle/IntelliWatch/Monitor/USER
```

– in *iwmon.ini* (UNIX)

```
[USER]
[relevant field]
```

The only such field created by the Setup is *Latest Rip Date*. To reference this field--or any custom fields you may have added--enter it in the Trigger's "Field (Date/Time or Bytes)" text box (in square brackets [ ]).

Example:

To receive a copy of the resulting NSD*.log/Notes.rip file when crashes occur, create a File Trigger that checks for a date change in the file, then sends it to the Admin.

The combo boxes in question *are* enabled when Size is the condition type, since units of size must be specific.

Trigger: File Trigger

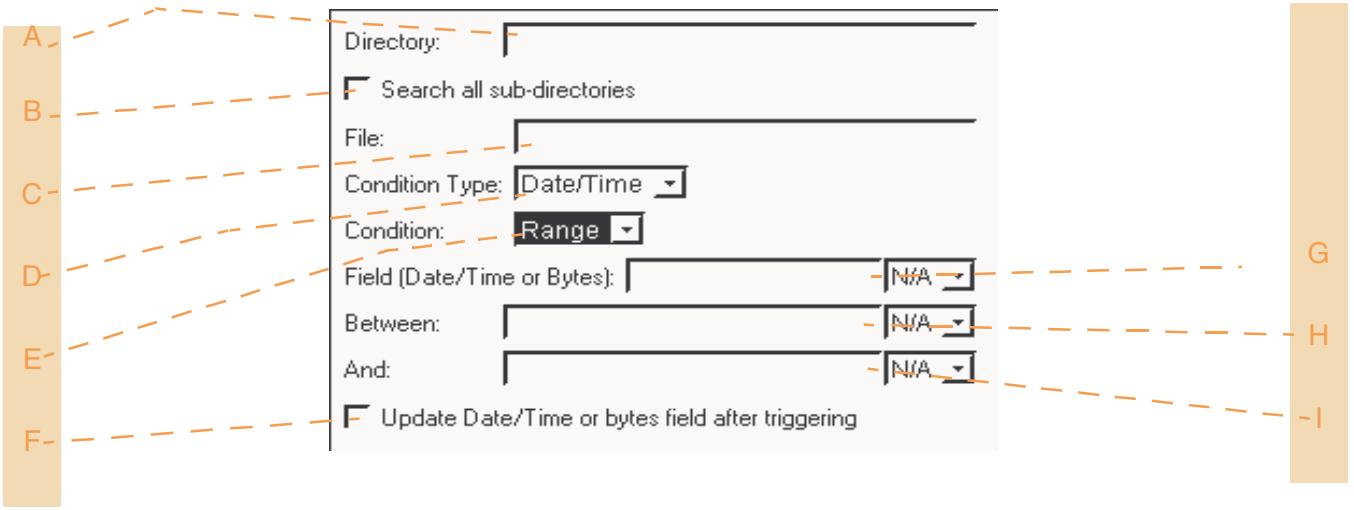

A: directory to search for file to be monitored
B: if selected, search all subdirectories
C: file(s) to monitor (wildcards permitted)
D: condition type (Date/Time or Bytes)
E: condition operator governing file check. Click on combo box arrow to access list (double-click icon for options).
F: if selected, updates Date/Time/Size information each time Trigger fires
G: field to monitor (for units, double-click icon)
H: lower limit (when condition operator is Range). End point values are included. (For units, double-click icon at G.)
I: upper limit (when condition operator is Range). End point values are included. (For units, double-click icon at G.)

Trigger Type: Replication Integrity Trigger

Basics ...

Checks for:

- database with differences after replication with specified server (wildcards are not supported)

In pre-27.33 maintenance releases of IntelliWatch, selective replication formulas could be taken into account, but no other document-selection filters were available.

Starting in 27.33, selection can be based EITHER on replication formulas, or on document types (for available types, see graphic below).

Selecting the replication-formula option overrides selection by document type. When you choose Select Document Types, All Document Types *is the default; check the* Select Document Types *radio button to limit the Trigger to only some types.*

Putting it into practice ...

Example:

Tree Technologies, Ltd. discovered that even when replication appeared to have proceeded normally, errors had sometimes occurred that left databases out of sync.

Admins used Replication Integrity Triggers to check for successful replication, and to notify them whenever problems occurred.

In addition to determining that replication of given databases had been unsuccessful, and notifying them of that fact via e-mail, Replication Integrity Triggers were configured to invoke an IntelliWatch Server Console Command (replica), to retry replication of the relevant database.

Trigger: Replication Integrity Trigger

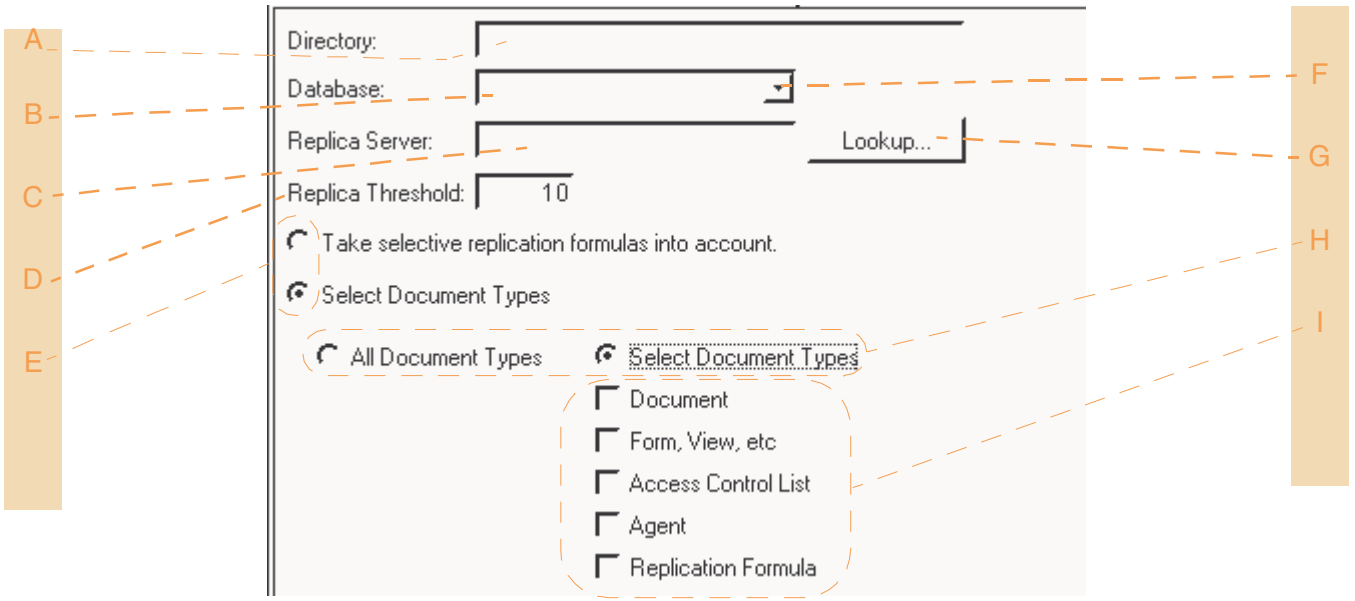

- A: full path to directory containing database(s). Leave blank (or enter [Directory] for the Notes data directory.
- B: database to examine
- C: select Replica Server
- D: set threshold for minimum number of differences required for Trigger to fire
- E: elect EITHER to use replication formulas, or to limit the check by document types
- F: select database using combo box, or type in name (field uses auto-complete for known databases)
- G: button launches Server Selection dialog
- H: choose EITHER all types, or use check boxes to limit Trigger to certain types
- I: document types available as selection criteria (multiple selection allowed)

Trigger Type: Replication Readiness Trigger

Basics ...

Checks for:

- Conditions that adversely affect replication, such as:
  - ACL conflicts
  - improper Replication settings
  - problems with connection documents

Putting it into practice ...

Example:

You have critical Notes applications that need to be kept current across your environment. Occasionally, however, scheduled replication is turned off while system maintenance is being performed.

Use this Trigger type to check the replication settings of these critical databases, and notify the relevant Admin(s) in the event a problem is encountered.

(All necessary access rights must be in place for all databases involved.)

Trigger: Replication Readiness Trigger

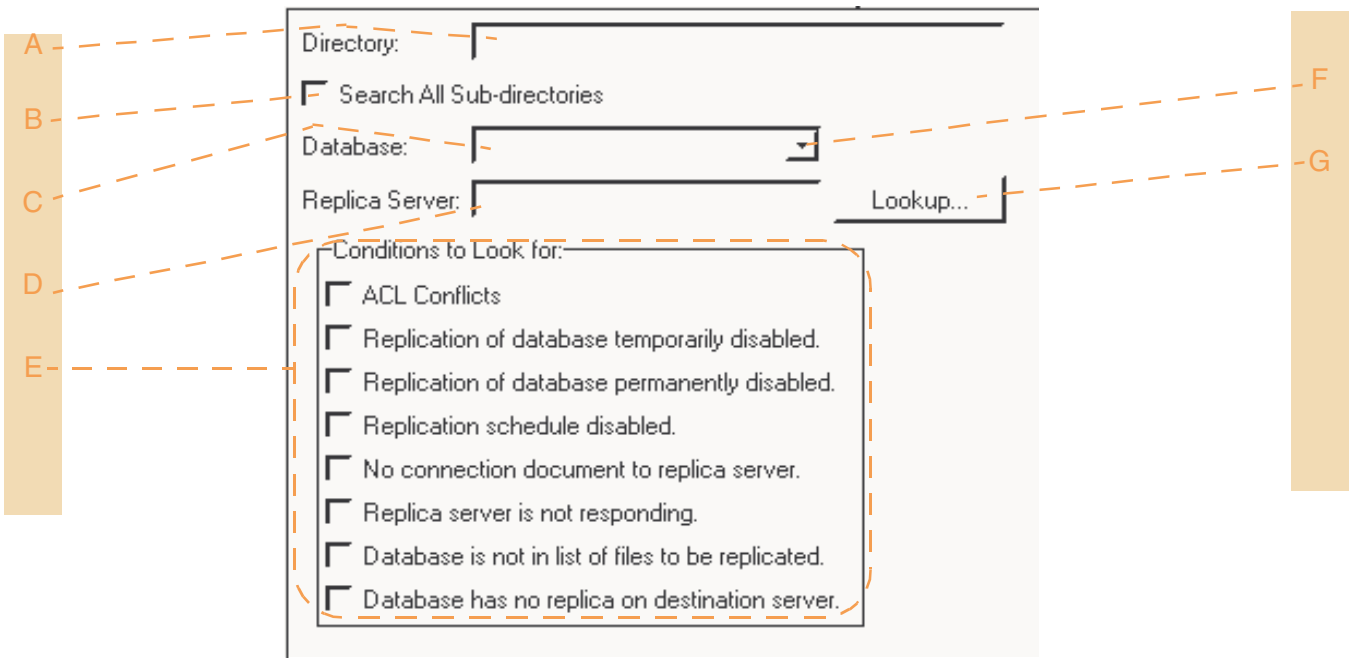

A: full path to directory containing database(s). Leave blank, or enter [Directory] for the Notes data directory.
B: check to search sub-directories
C: select database to examine
D: select Replica Server
E: select all desired conditions affecting replication
F: use combo box to select database
G: button launches Server Selection dialog

Trigger Type: Statistic Trigger

Basics ...

Checks for:

- Domino server statistics
- IntelliWatch Performance Manager statistics

Check against user-specified levels to identify problems, as well as fine-tune your environment.

PM Statistic Names and IntelliWatch Monitor:

There is a difference between the Name given to PM statistics at the Domino server console, and how those same statistics are listed in the IntelliWatch databases.

In both **iwpmstat.nsf**, where PM statistic definitions are stored, and in **iwpmstats.nsf**, where PM statistics are collected, the Statistic Name is identical to that created/displayed at the Pinnacle Console.

At the Domino server console, on the other hand, **"iwstats."** is pre-pended to the Name, to help you differentiate IntelliWatch statistics from Notes statistics. (To show all IntelliWatch statistics, type **sh st iwstats**.)

When manually adding PM statistics to the list displayed in a Trigger, therefore, you must add the **iwstats** prefix, or the Trigger will not function correctly.

Putting it into practice ...

Example:

The Notes statistic Server.Sessions.Dropped is unacceptably high on certain servers in your environment, leading to slow server performance.

This statistic reflects the number of users who gave up trying to connect to a server, usually by restarting or pressing CTRL+BREAK.

Use a Statistic Trigger to monitor for a given number/range of dropped sessions. If it happens regularly, it could signal the need for a hardware upgrade, database cleanup, more memory, or another server.

Statistic Triggers can *compare non-numeric values. Note, however, that the comparison is then lexicographic, or character-by-character. The number of places in each value is critical.*

Comparing a Server.ElapsedTime value of

- *6 days 05:23:43*

to Trigger search values of

- *25 days 00:00:00* (too many places) or
- *12:00:00* (too few places)

will behave unpredictably.

Trigger: Statistic Trigger

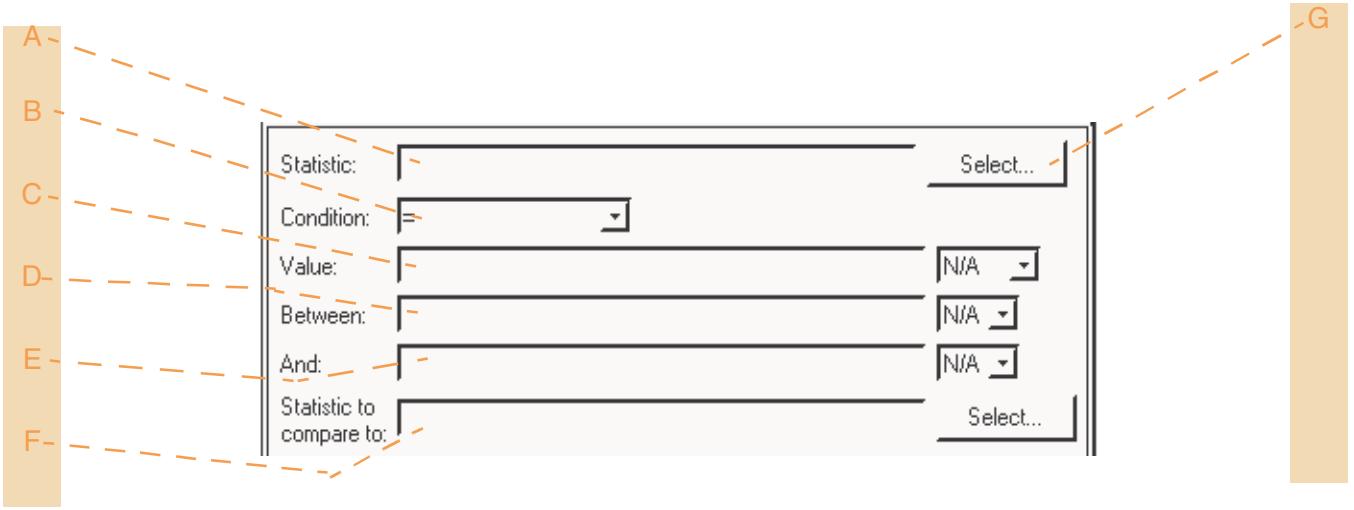

A: select Statistic to be monitored by clicking button at G

B: select condition. Click on arrow of combo box to access list (double-click icon for options).

C: value required for Trigger to fire. Click on the arrow of the combo box to access the list of units (double-click icon for options).

D: lower limit (when condition operator is Range). End point values are included. (Units at C.)

E: upper limit (when condition operator is Range). End point values are included. (Units at C.)

F: second Statistic to be checked as a percentage of the first Statistic (in field A). This field is enabled *only* when % is selected as the unit of Value.

G: button launches Select Server dialog

Trigger Type: User Trigger

Basics ...

Checks for:

- exit status of programs
  - Use with in-house server monitoring utilities, for example.

Monitor launches the user-supplied utility or batch file during the monitoring cycle, and, based on the exit status returned from the utility, executes any IntelliWatch Commands you specify.

Program must meet two requirements:

- Must complete running in a reasonable period of time (approx. 2 minutes), because IntelliWatch monitoring cycle is suspended until program exits.
- Must exit with return code which signals exit status. Monitor uses return code to determine if Commands are to be executed.

Putting it into practice ...

Example:

Ace ISP has an in-house utility, CHECK_SVRERROR, that checks for a particular error condition, and returns a value of 1 when it's detected (a return code of 0 signals NO_ERROR).

They created a User Trigger and specified CHECK_SVRERROR as the program to run. Whenever the Trigger detects a return value of 1, Monitor sends a page, and launches CORRECT_SVRERROR.bat to remedy the issue.

Trigger: User Trigger

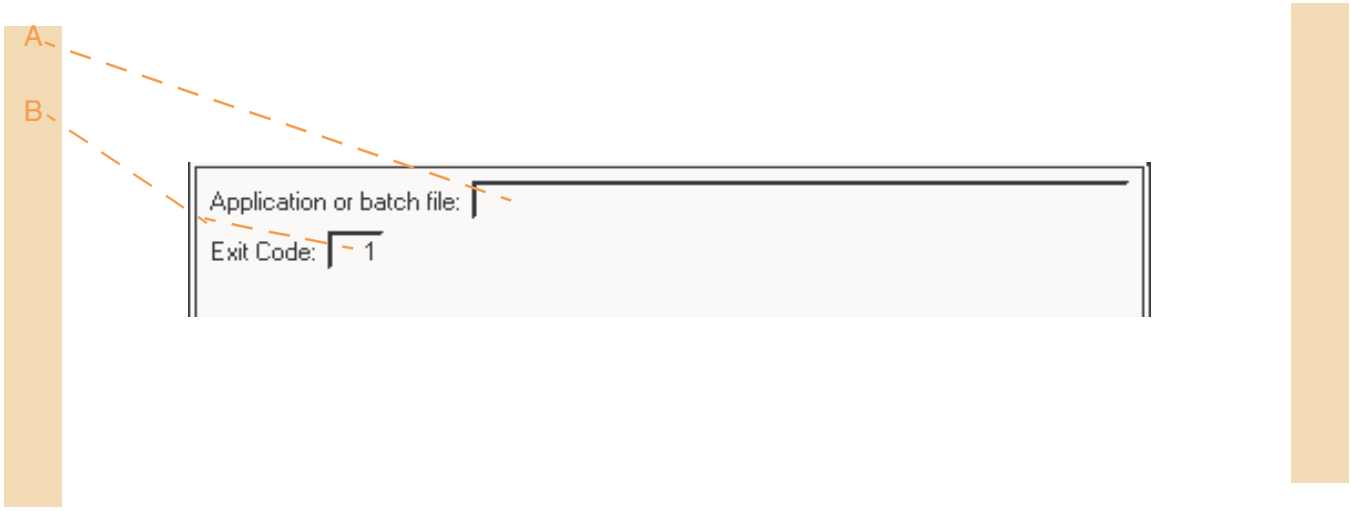

A: name of user-supplied program to run
B: exit value or return code for which to test

Trigger Type: White Space Trigger

Basics ...

Checks for:

- amount of White Space in databases

Excessive White Space negatively impacts performance and wastes valuable disk space. Therefore, you should regularly check for it and remove it.

When a White Space Trigger indicates that a predefined threshold has been reached, you can automatically initiate compaction to improve performance and reclaim disk space.

Wildcards are supported (for example, *.nsf). *However,* if wildcards are used, a Whitespace Trigger ignores **names.nsf**, **mail.box**, **log.nsf** and other databases which are usually open.

=, <>, >, < and **range** operators apply only to numeric fields.

Putting it into practice ...

Example:

Just as with other databases, excessive White Space in Log.nsf can adversely affect performance. Unlike most other databases, however, to compact Log.nsf, you must first stop the Notes server.

To accomplish these tasks using IntelliWatch Monitor:

- Create a White Space Trigger to monitor Log.nsf.
- Set up a Recycle Command to be issued to the Server whose Log.nsf needs to be compacted.
- Associate a batch file with the Recycle Command that runs Compact on Log.nsf before restarting the server.

Trigger: White Space Trigger

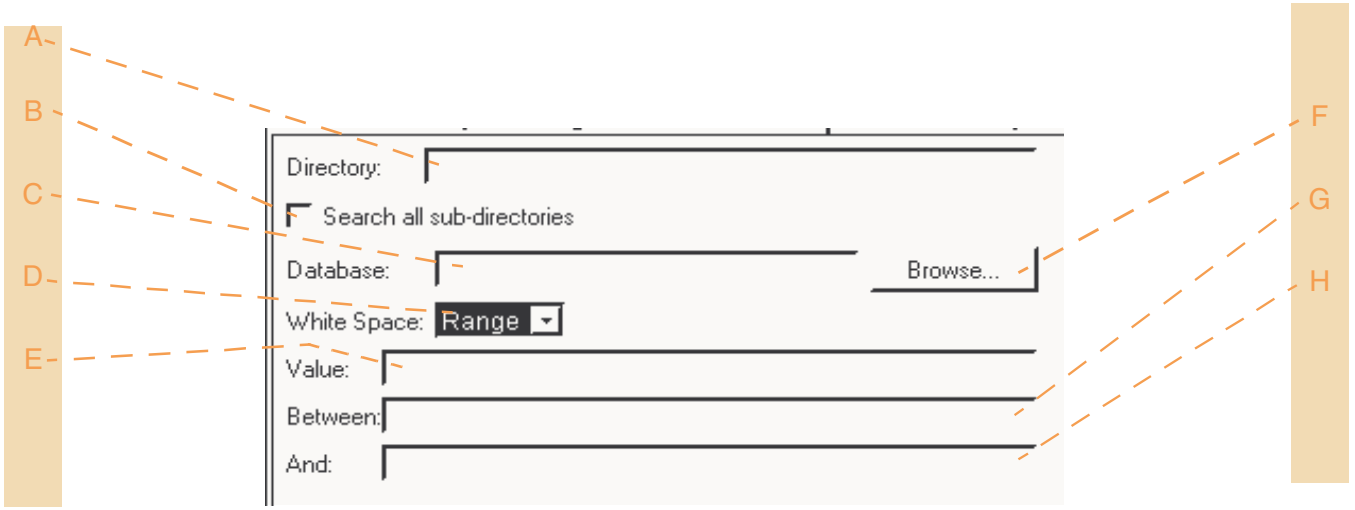

A: location of database(s) to monitor
B: if selected, search all subdirectories
C: database(s) Trigger should monitor. Wildcards are supported (for example, *.nsf).
D: operator type for condition statement (double-click icon for options)
E: value to search for, as percentage of file size (disabled when Range operator selected)
F: button launches Select Database dialog
G: lower limit (when condition operator is Range). End point values are included.
H: upper limit (when condition operator is Range). End point values are included.

Command Type: ACL Changer

Basics ...

What it does:

- adds User/Group to ACL of database
- changes access privileges for existing User/Group

Putting it into practice ...

Example:

The Global Services Company's Notes team created a new group called UsrMailAdmins to manage users' mail databases across their environment.

To simplify adding this group to the ACL of user mail files, they invoked the ACL Changer Command.

Since the affected mail files were the only Notes databases residing in the directory *Company\Users*, Admins configured the Trigger invoking the Command to add the new group to the ACL of *Company\Users*.nsf*.

Command: ACL Changer

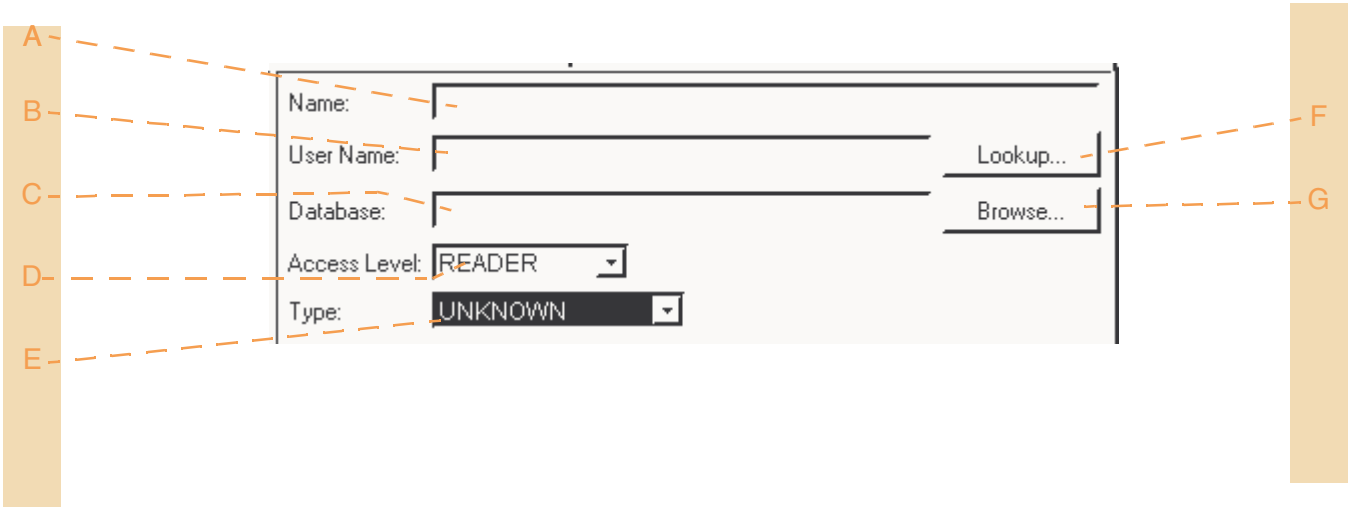

A: name of Command
B: user/group for which to change/add access
C: name of database(s) whose ACL are to be changed.
D: select new level of access. Click on arrow of combo box to access list (double-click icon for examples).
E: select type of owner of database. Click on arrow of combo box to access list (double-click icon for examples).
F: button launches NAB
G: button launches Select Database dialog

Command Type: IW Alert

Basics ...

What it does:

- sends an alert to the Advanced ServerWatch message hub

Messages are assigned one of three priorities:

- 1=green (a status message, such as "IntelliWatch Monitor finished initializing")
- 2=yellow (warns of a possible problem, such as a higher than normal response time)
- 3=red (error condition)

Putting it into practice ...

Why use IWAlerts in addition to e-mail and paging?

Three features of IWAlerts make them an ideal messaging tool:

- IWAlerts are displayed by originating server
  - The ASW database (**iwasw.nsf**) sorts incoming alerts by the monitored server that sent them. Manually categorizing notifications (as with mail messages/pages) is unnecessary. The ASW archive database allows for convenient storage of IWAlerts for the longer term.

Depending on the number of messages present for a given server, adjust the settings at **File > Preferences** to increase the number displayed (if you do not see all IWAlerts for a given server).

- Color-coded for severity
- Message pairing
  - By means of the message-pairing feature, Admins can see not just when a problem occurs, but when the situation has been corrected. (For more information, see *n on page 169*.)

Example:

IWAlerts are a particularly useful notification tool when you have a staffed monitoring center. Clicking on the icon for a given server causes the Console to display applicable IWAlerts. Depending on the nature of the message, corrective actions can be initiated, or the proper people/department can be notified.

(The ASW Console gets its information from **iwasw.nsf**. on the ASW Hub. To limit the ability of monitoring staff to make changes in that database, use the ACL.)

Command: IWAlert

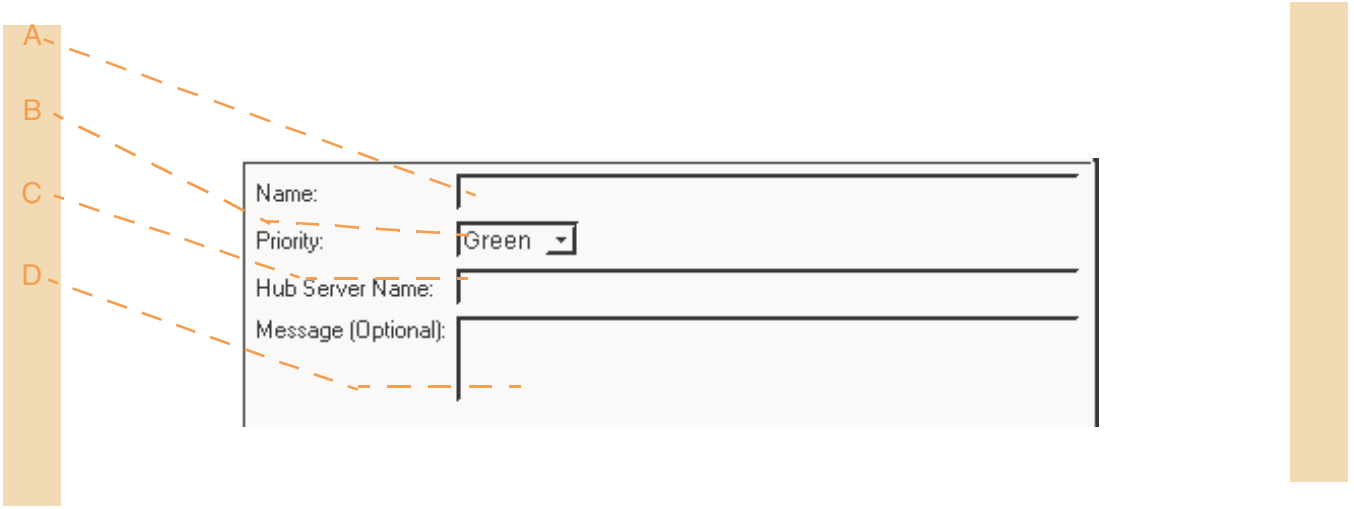

A: name of Command
B: message priority. Click on arrow of combo box to access list (green, yellow or red)
C: ASW Hub Server
D: optional message to be appended to message sent by Trigger invoking Command

Command Type: IWNTLog

Basics ...

What it does:

- sends a message to the NT Events Log

The server parameter refers NOT to the Notes server, but to the destination NT system for the Event. (Leave blank for the Local NT Event Log.)

Putting it into practice ...

Types and categories:

- Types
  - Audit Failure
  - Audit Success
  - Error
  - Information
  - Warning
- Categories
  - General
  - Communication
  - Corruption
  - Database
  - Database Corruption
  - Document Corruption
  - Full Text Corruption
  - Mail
  - Replication
  - Resource
  - Server
  - Server Access
  - Template

Command: IWNTLog

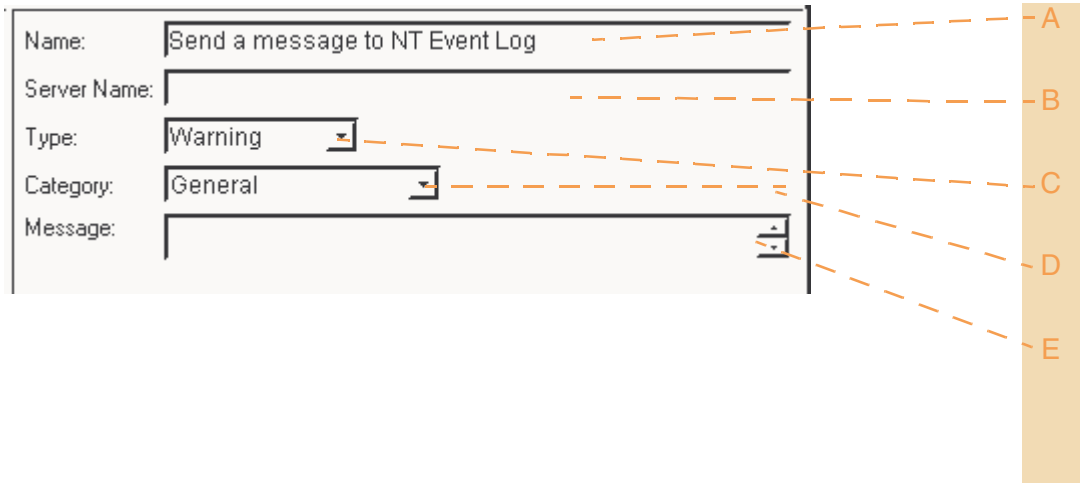

A: name of Command
B: server whose NT Event Log is to receive message
C: NT Event type
D: NT Event category
E: message to be sent with Event

Command Type: IWMail

Basics ...

What it does:

- sends e-mail via Notes to the designated user(s)

IWMail Commands are used to send e-mail to any Notes mail system. Messages are passed to **iwmail.nsf**, which hands them off to **mail.box** for routing.

Both short user names (user_name) or fully qualified names (user_name@DOMAIN) can be used for mail recipients. Separate multiple names with commas.

The keyword <MAILUSER> can be used in TO and CC fields.

Either <FILE> or <DATABASE> keywords can be used in Import Files and Attach Files fields, if called from Triggers supporting those keywords.

Putting it into practice ...

Message field:

Any message entered here is appended to the message contained in the Trigger invoking the Command.

Command: IWMail

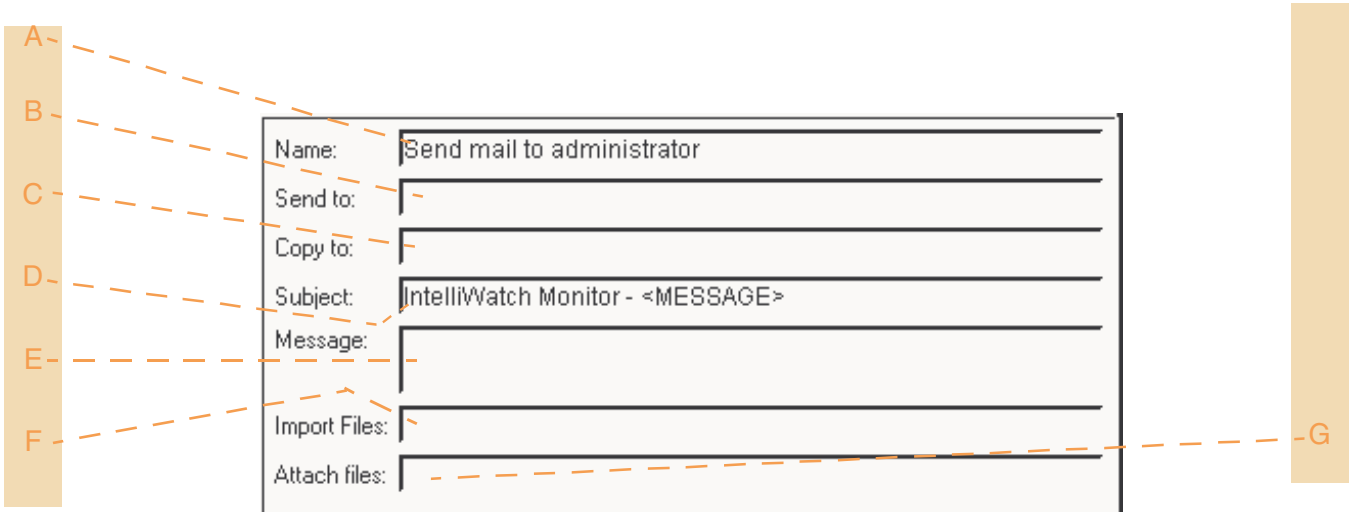

A: name of Command
B: comma-delimited list of mail recipients
C: comma-delimited list of mail recipients to be copied
D: subject of message (enter keyword <MESSAGE> to put first 220 characters of message into subject field)
E: message body (any message entered here is appended to message in Trigger)
F: comma-delimited list of files to be imported into message
G: comma-delimited list of files to be attached to message

Command Type: Kill Process

Basics ...

What it does:

- kills a named process

Putting it into practice ...

Restrictions:

On UNIX, the User invoking the Command must have the proper permissions for the process in question.

On NT, similar restrictions depend on any network security protocols that are in place.

Command: Kill Process

A: name of Command
B: name of process to kill

Command Type: Move File

Basics ...

What it does:

- moves file or directory from one drive to another, or from one location to another on the same drive

Putting it into practice ...

Example:

Use this Command type to free up disk space (to allow the swap file to grow).

If the Trigger invoking this Command allows use of the <FILE> and <DATABASE> keywords, enter them in the Source File field.

Wildcards are supported (for example, *.nsf).

Command: Move File

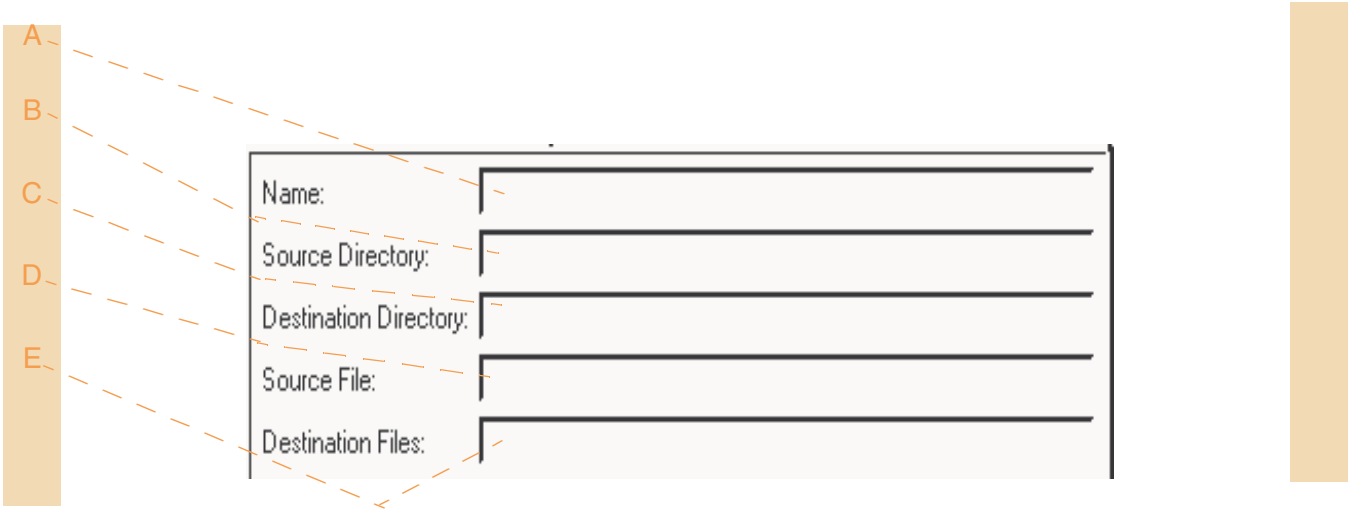

A: name of Command
B: path to files you want to move
C: path to destination directory
D: comma-delimited list of files to be moved
E: comma-delimited list of destination files

Command Type: Move/Remove Document

Basics ...

What it does:

- moves documents to new Notes database, or removes them entirely

This Command type allows you simultaneously to control the size of Notes databases, and to archive documents you want to preserve.

If a destination database is specified (see field D on the following page), documents are moved; if the destination database field is left blank, documents are deleted.

Putting it into practice ...

Example:

Use this Command type to archive documents and/or control the growth of databases.

If the Move/Remove Command is called from a Replication Integrity or Document Time-out Trigger, the keyword <Document_ID> can be used in the *Document ID list* field (see letter C on the following page).

When a Document Time-out Trigger calls a Move/Remove Command that uses the keyword, any documents identified as meeting the time-out criterion will be moved/removed (depending on the value of the Destination Database field).

Notes documents identified as having exceeded the configured threshold of a Document Time-out Trigger. for example, can be moved to another database—call it ArchiveOld.nsf—or deleted.

Command: Move/Remove Document

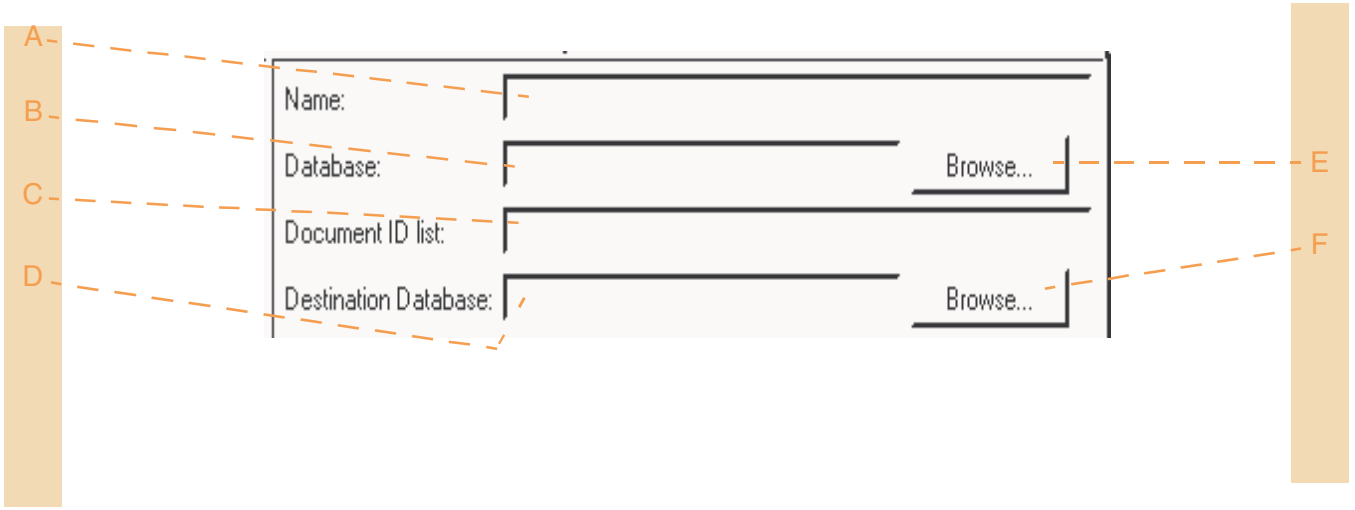

A: name of Command
B: database containing document to be moved/removed
C: ID of document to be moved or removed (may be comma-delimited list)
D: database to which to move document (if this field is left blank, document is deleted.)
E: button launches Select Database dialog
F: button launches Select Database dialog

Command Type: New Replica

Basics ...

What it does:

- creates a new replica of a Notes database

Putting it into practice ...

Example:

If a Trigger detects corruption in a database, you may want to invoke this Command type to create a new replica on the server in question.

(If the New Replica Command is called from a Database Trigger, use the keyword7 <DATABASE>, <NOTES_DB1> or <NOTES_DB2> in the Database List field.)

Command: New Replica

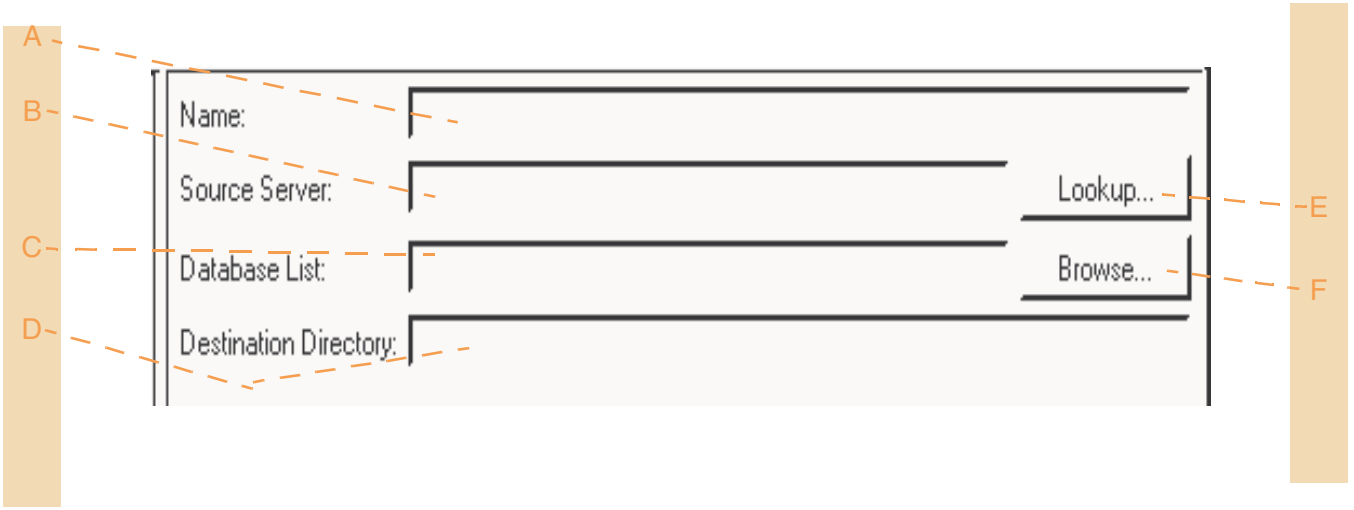

A: name of Command
B: name of server on which databases to be replicated are located
C: list of databases to replicate
D: directory in which to put new replica (leave blank for the Notes data directory)
E: button launches Select Server dialog
F: button launches Select Database dialog

Command Type: Pager

Basics ...

What it does:

- routes text message to alphanumeric pager

Pages logged to **iwpage.log** (in IntelliWatch Logs directory—specified in the NT registry on the Message Center Gateway). Delete the file if it becomes too large; it will be recreated as needed.

Putting it into practice ...

'Local' Setting:
A Comm Port, Dial Type or Port Speed set to "Local" causes the server to use paging settings where the command is actually run.

Similarly, whenever other fields are left blank, the paging server will use paging settings on the system on which the Command is executed (and not the default settings on the Paging Server itself).

Name:
The Name field is not populated by the Paging Server, and each Pager Command must be given a unique name before it can be saved.

Message:
Field limit of 500 characters may be influenced by service provider. Any message entered here is appended to the message in the Trigger.

Initialization string:
Leave blank, unless experiencing problems (refer to modem manual for an initialization string that sets modem to BELL 103 compatibility).

Tone or pulse:
Verify with your telecom staff whether TONE or PULSE dialing is supported at your location)

Comm port:
Other programs running on the system (such as Notes) must not have previously allocated this port.

Command: Pager

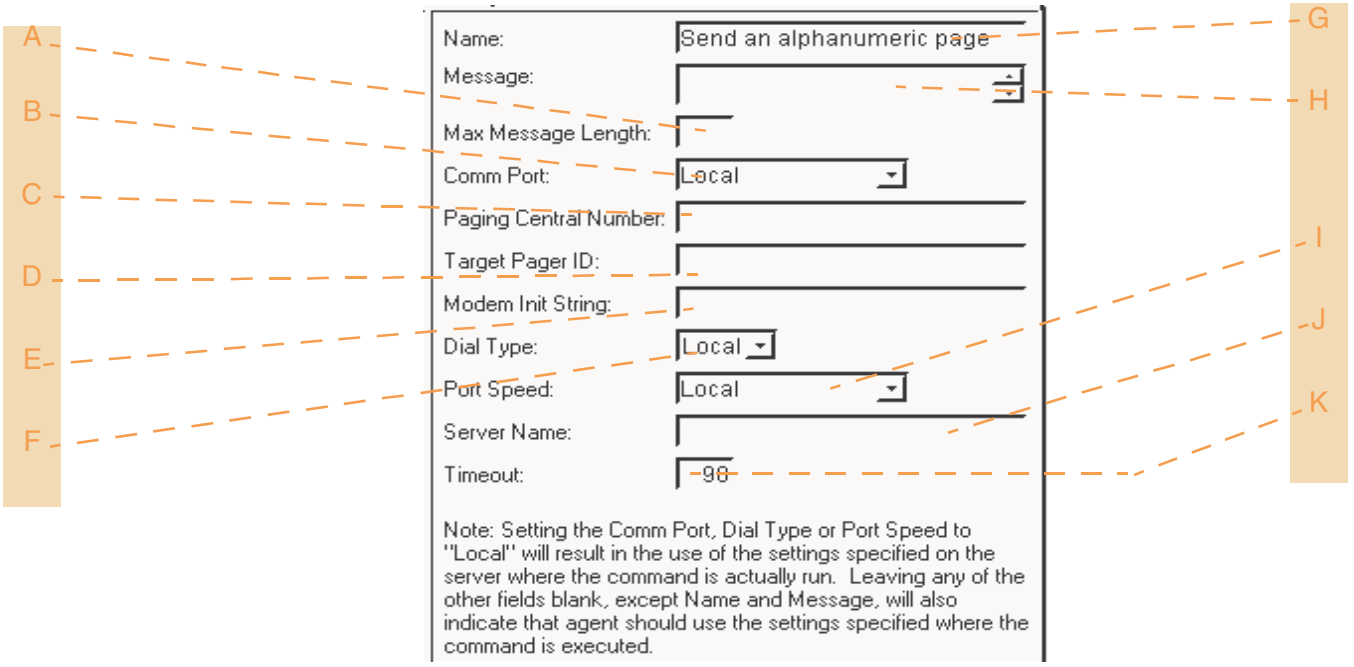

A: maximum number of characters allowed
B: port to which modem is attached
C: paging central's modem number
D: unique ID for your alphanumeric pager
E: modem initialization string (leave blank unless modem not initializing properly)
F: tone or pulse dialing (for value of 'Local', see preceding page)
G: name of Command
H: message to be sent
I: speed at which modem and serial port communicate
J: TCP/IP Hostname of Paging Server
K: how long to wait (in seconds) before terminating pager connection

Command Type: Reboot

Basics ...

What it does:

- reboots the system at the OS level

Graceful vs Immediate:

- graceful

– The Notes server (including any associated tasks) is stopped before the system is rebooted.

- immediate

– Reboots the system without stopping the Notes server.

Putting it into practice ...

One effective use of the Reboot Command is in batch files that automate your maintenance cycles.

Example:

Certain Notes databases (such as **log.nsf**) either cannot be compacted while the server is running (R4), or usually are not (R5), due to performance consequences. Doing maintenance on these databases therefore means bringing down the Domino server, at which time COMPACT, UPDALL, and like programs, are run. To automate this using IntelliWatch:

- create a Trigger that you know will evaluate to True
  - – Look for a file you know exists, for example.
- launch a first batch file that
  - – issues a 'quit' to the server, followed by
  - – the relevant Notes server commands, then
  - – **iwrecycl** (see *page 138*).

(Note, however, that one of the invoked tasks may hang. The batch file then never gets as far as the Recycle Command, and the server does not come back up.)

- launch a second batch file as 'insurance'

```
iwsleep [maintenance cycle in seconds, plus a short wait]
iwreboot [REBOOT_IMMEDIATE or REBOOT_GRACEFUL]
```

(For usage details, type iwreboot at a DOS prompt.)

Now, even if a server task hangs, the machine will be rebooted by IntelliWatch to bring it back on-line.

(Note: The batch files must be installed locally on all servers running the Trigger, and should be in a directory included in the path, such as \Candle\IntelliWatch\Common.)

Command: Reboot

A: name of Command
B: Immediate or Graceful (for usage, see preceding page)

Command Type: Recycle

Basics ...

What it does:

- recycles Notes server and associated tasks

In order to restart the Notes server successfully, any program which has accessed Notes must be stopped. This is handled internally by the program, without your having to enter program names.

Putting it into practice ...

Notification levels:

- 'None': no notification
- 'Limited': notifies server has been restarted
- 'Verbose': notifies at all significant stages of recycling.

Use the Verbose setting only when instructed to do so by IntelliWatch Customer Service personnel, as it generates a considerable number of notifications.

Example:

In addition to responding to server crashes, Recycle can also be used to kick off scheduled maintenance.

Simply create a Program Document in Notes to launch Recycle at the appropriate time.

Command: Recycle

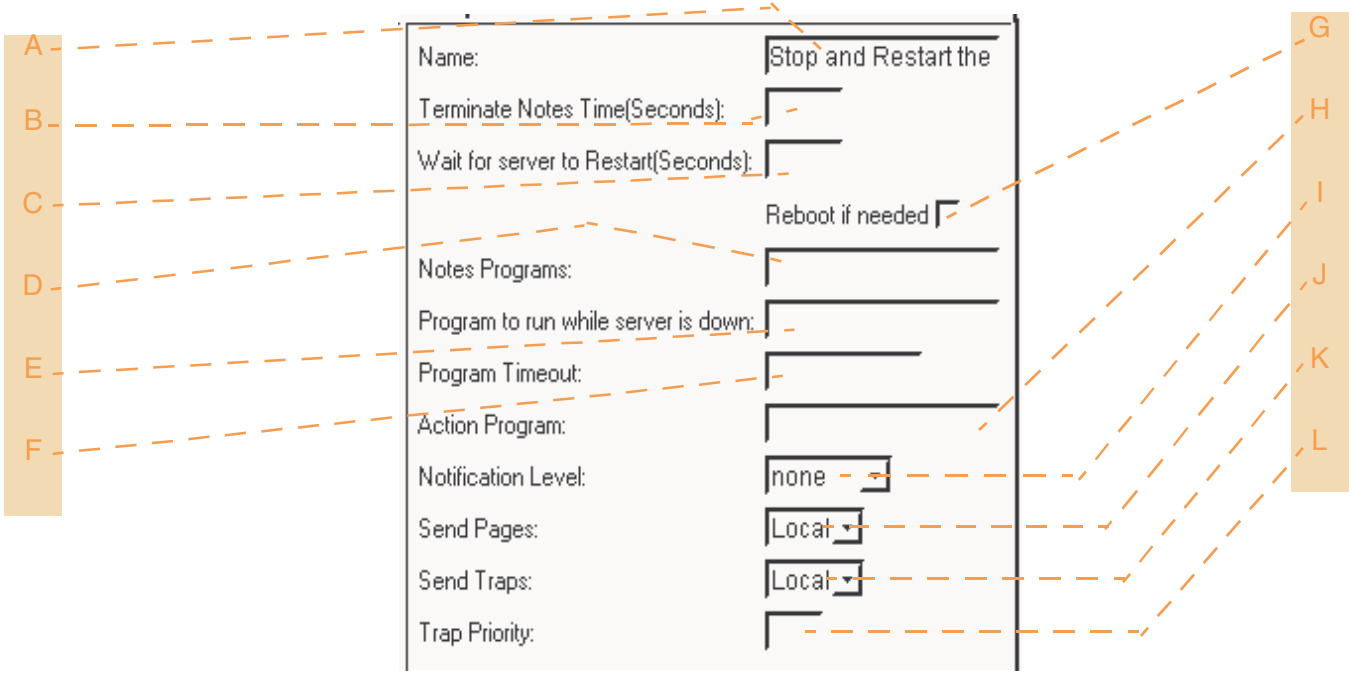

- A: name of Command
- B: time (in seconds) to wait for server to terminate
- C: time (in seconds) to wait for server to restart
- D: list of programs not part of standard installation
- E: program to run while server down
- F: time (in seconds) to wait for batch file to terminate before restarting Domino server
- G: reboot if needed (when selected)
- H: program to run each time notification sent
- I: notification level (none, limited, or verbose) For usage, see previous page.
- J: if selected, send page to notify (local, no, or yes)
- K: if selected, send SNMP trap to notify (local, no, or yes)
- L: priority of SNMP traps sent as notifications

Command Type: Restart Add-in

Basics ...

What it does:

- restarts a specified Notes Add-in task

Putting it into practice ...

Example:

Consider using this Command type to invoke third-party Notes add-ins. Please consult the vendor in question for required/optional command-line arguments, as well as any other information necessary to properly configure the add-in.

Command: Restart Add-in

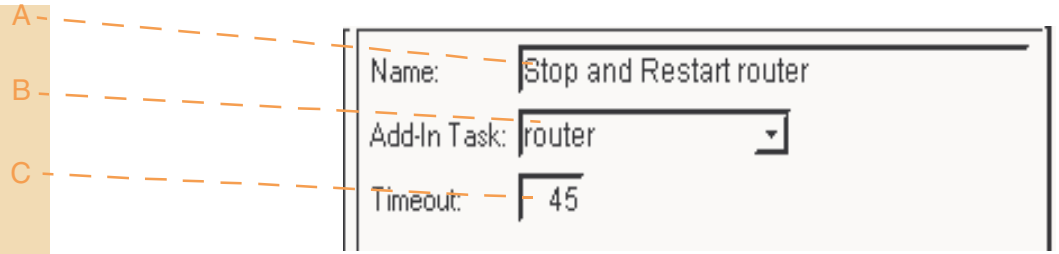

A: name of Command
B: process name of Notes Add-in to stop and restart (double-click icon for examples)
C: time (in seconds) to wait for Add-in task to terminate

Command Type: Run Agent

Basics ...

What it does:

- runs a specified agent against a database

Putting it into practice ...

Example:

You have intermittent problems with dead mail on certain servers. An Agent has been created to deal with these messages, but you would rather not have to invoke it on a schedule, preferring to launch it only when necessary.

A Statistic Trigger can check the value of Mail.Dead on the servers in question, then launch the Notes Agent when the value of the statistic is above the configured threshold.

Command: Run Agent

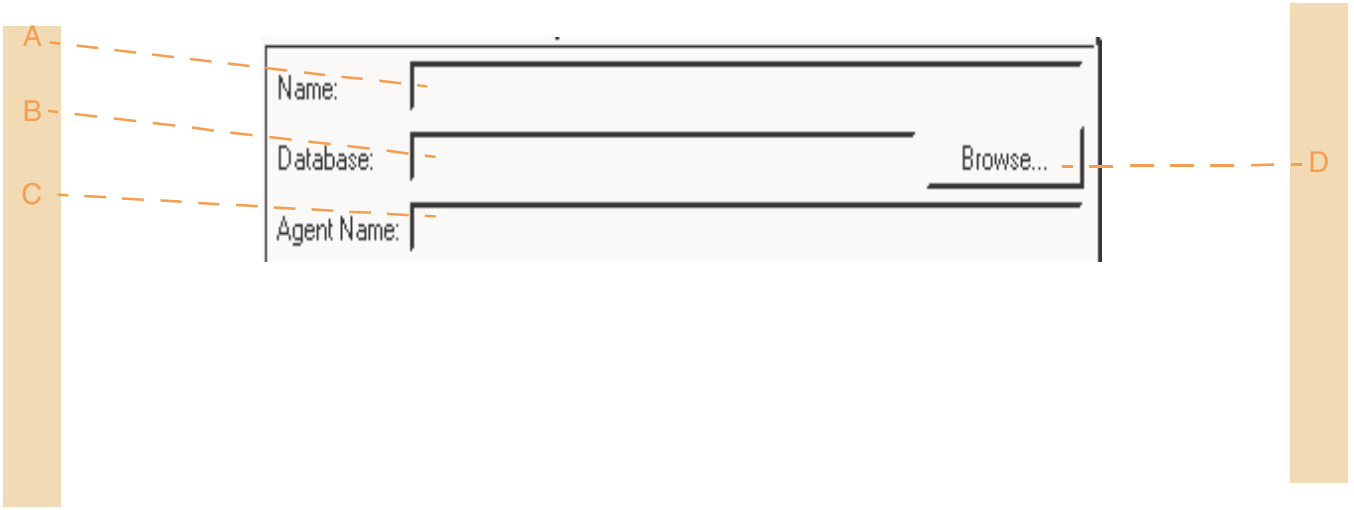

A: name of Command
B: database containing agent and on which agent runs
C: name of agent to run
D: button launches Select Database dialog

Command Type: Server Console

Basics ...

What it does:

- executes a Domino Server console command, such as "load router"
  - issues commands to local *and remote* servers

Click on the Parameter button for help on which arguments can be used with a given Command.

Putting it into practice ...

Example:

You prefer to use ReplicaServer for the scheduled replication of **XYZ.nsf**. However, ReplicaServer is occasionally unavailable at the required time.

Use a Replication Readiness Trigger to check for the availability of ReplicaServer. If a problem is detected, have the Trigger issue a Server Console Command (that invokes "load replica", and includes a reference to an alternate replication server).

Either the <FILE> or <DATABASE> keyword can be used in the parameter field. (This causes the database referenced by the Trigger to be included as an argument in the Command to be executed.)

Command: Server Console

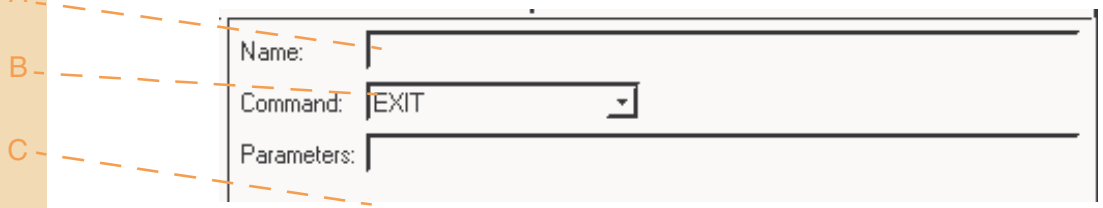

A: name of Command
B: drop-down used to select command (double-click icon for examples)
C: parameters/arguments

Command Type: Sleep

Basics ...

What it does:

- causes IntelliWatch to sleep for X number of seconds

Putting it into practice ...

Example:

When you create batch files that you want to call with IntelliWatch Start Program Commands, sections of those batch files may require a measurable amount of time to complete.

Problems can result if subsequent commands in the batch file are executed without allowing time for previously issued commands to finish.

IWSleep was created to address this need.

Command: Sleep

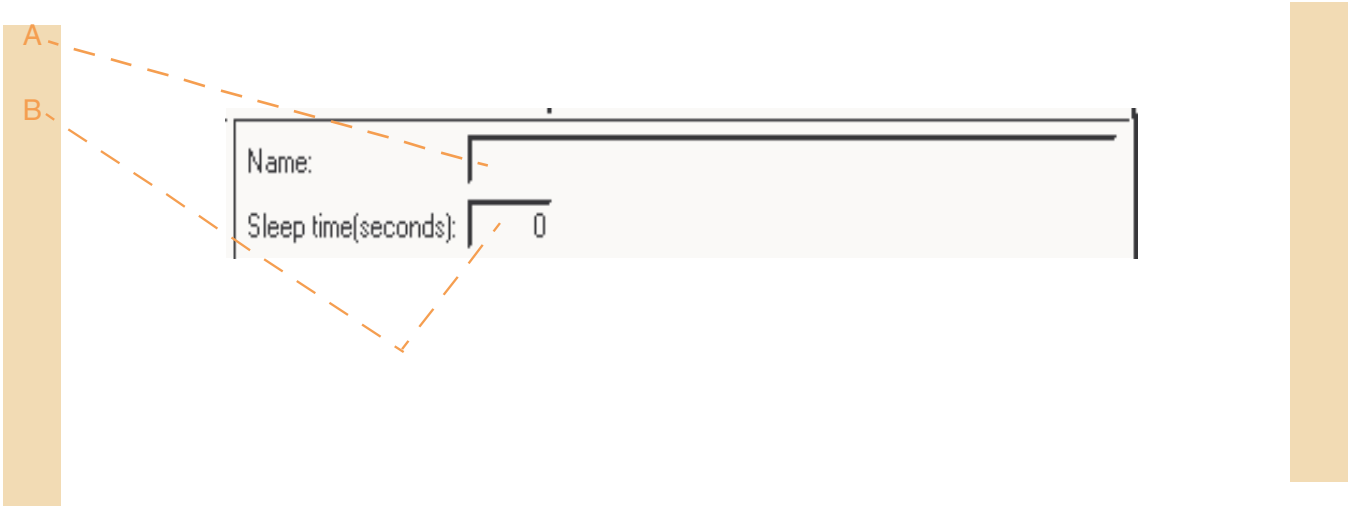

A: name of Command
B: how long to sleep (in seconds) before proceeding

Command Type: SNMP Trap

Basics ...

What it does:

- sends an SNMP trap to the defined SNMP consoles

Putting it into practice ...

Hostname:
May be either a TCP/IP address (124.34.216.1, for example) or a HOSTNAME environment variable specified in config.sys (such as SET HOSTNAME=[your hostname]).

Community name:
Specify a community name, or leave blank for the default.

Priority:
A priority number is sent to the SNMP Network Manager as part of the trap, and can be used for additional processing in a script. If no priority is specified, the system default is used. You may be able to map this field to a severity level at your SNMP Network Manager.

Command: Send SNMP trap

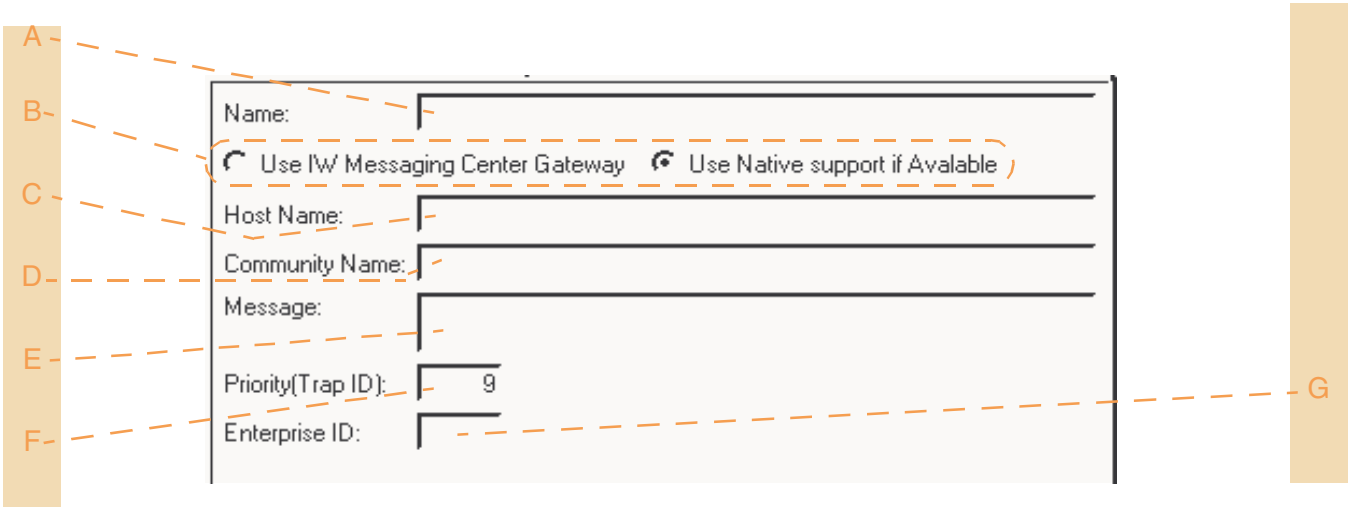

A: name of Command
B: select IW Messaging Center Gateway or native support (if available)
C: hostname of machine on which Monitor is installed
D: authentication mechanism used by SNMP to verify requests
E: any message entered here is appended to message sent by Trigger
F: set field to any number
G: enterprise ID of trap (seldom needed)

Command Type: Start Program

Basics ...

What it does:

- executes (system) commands

Use any IntelliWatch keyword in this field, enclosed in <> symbols (<keyword>).

The Program Arguments field must respect the syntax of any batch file you are invoking, including things like flags and keywords.

For instance, to send an SNMP trap to the Message Center Gateway that includes the text of any message being sent by the Trigger, configure it as follows:

- Batch file or program name: iwevent
- Program arguments: SNMP /M:<MESSAGE>
  - Do NOT use quotes around the last argument, or the Command will fail.

Putting it into practice ...

Notification command:

If a program is specified as a notification program (see letter F on the following page), Monitor runs it at each escalation of a Trigger. If this parameter is set to false (isn't selected), a notification program runs only once, on the First Occurrence (unless specified at each escalation).

Command: Start Program

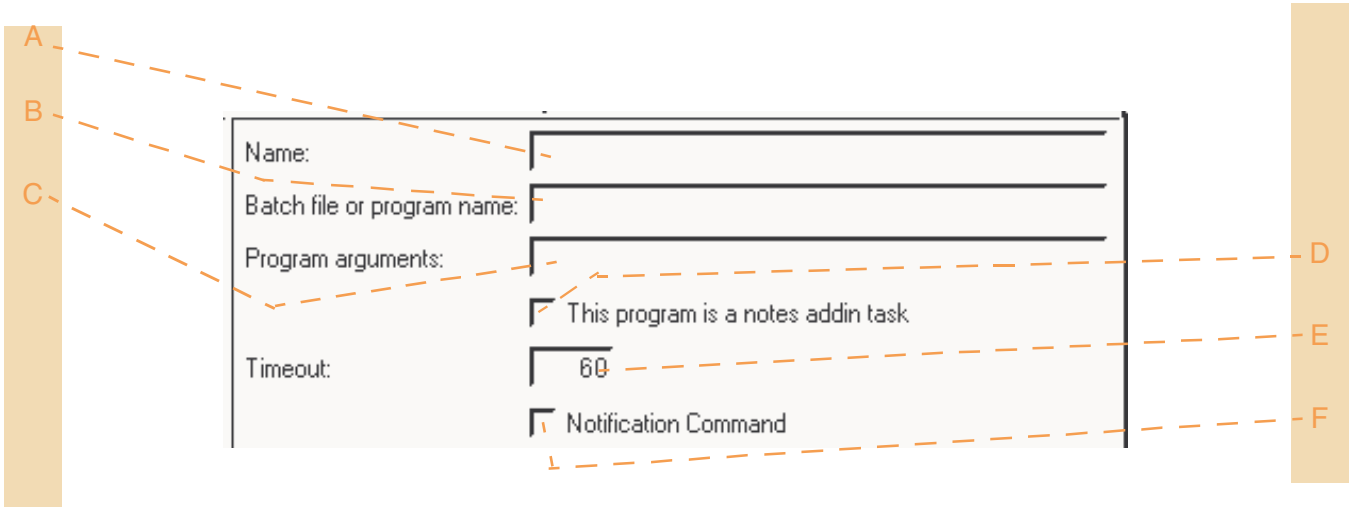

A: name of Command
B: name of program or batch file to launch
C: list of options and switches to set for program
D: if selected, program is a Notes add-in task
E: time (in seconds) to wait for program to complete before executing other Commands in Trigger
F: if selected, notification command

Command Type: TEC Event

Basics ...

What it does:

- sends IntelliWatch TEC Event Commands to Tivoli

Putting it into practice ...

Additional slots:

Specify additional slots using the syntax:

"<slotname>=<value>; ...; <slotname>=<value>;"

See your Tivoli documentation for additional information.

Command: TEC Event

A: name of Command
B: host name of machine running IntelliWatch Messaging Center Gateway. Machine must have TME support. Leave blank for default.
C: an event class defined in tecad_iw.baroc
D: severity of TEC Event. Click on arrow of combo box to access list (double-click icon for examples).
E: status of TEC Event. Click on arrow of combo box to change option (open or close).
F: any message included here is appended to message in invoking Trigger
G: space for supplementary comments
H: specify additional slots (for syntax, see previous page)

Chapter

Advanced ServerWatch

4

Advanced ServerWatch (ASW) supplies the tools you need to stay on top of the 'state of Domino' in your environment.

ASW accomplishes this by means of:

- an easy-to-read display of server connectivity
- three levels of error/status messages
- flexible Maintenance and Action profiles

Chapter Contents

4.1.0.0 OVERVIEW

Advanced ServerWatch provides real-time monitoring of remote-server availability, enabling you to see the state of your Domino server network at a glance.

The ASW Console (part of the Pinnacle Console) displays server connectivity status, as well as messages sent by IntelliWatch Monitor and the **iwalert** command-line utility.

Through user-created Action and Maintenance Profiles, tailor notification methods and maintenance cycles to the server's role in your environment.

4.1.1.0 Architecture

ASW consists of two components:

- ASW Server (Notes add-in task **iwasw**)
- ASW user interface

4.1.1.1 ASW Server

Runs on your ASW Hub servers *only*, and:

- checks server connectivity at a user-specified interval
- processes messages from **iwagent** *on the servers being monitored*

  Messages are stored in **iwasw.nsf** *on the ASW Hub*, located in the same folder as your Management Agents (R5 default /Lotus/Domino/Data/IntelliWatch).

4.1.1.2 ASW user interface

Accessible:

- via Internet Explorer
- through the stand-alone client

From the Console, view server connectivity status, configure the instances of ASW Server running on your various Hubs, and view IWAlerts by individual server.

No matter how small your network, however, you should set up at least two ASW Hubs.

4.1.1.3 Guidelines

- As far as possible, *ASW Hubs should monitor servers in their own geographical area.* Having a group of servers in Singapore monitored by a Hub in New York may create "false alarms." due to periodic connectivity problems.
- *Limit the number of servers reporting to a given ASW Hub.*

Having more than 50 servers reporting to a given ASW Hub is not supported.

- *Don't forget to monitor the ASW Hub itself.* If you have a single ASW Hub and it goes down *all the servers it's monitoring are "invisible" to ASW—until you restart it.*

4.1.1.4 Connecting to Pinnacle 99 Hubs

The Pinnacle Console has read-only access to Pinnacle 99 ASW Hubs.

4.1.2.0 Scenario 1:

ASW deployment in a seven-server environment

- ASW Server is installed on ASWHub1 and ASWHub2.
- ASWHub1 monitors all other servers, including ASWHub2 (during the installation of Monitor on these servers, they were configured to report to ASWHub1, which includes sending IWAlerts).
- ASWHub2 only monitors ASWHub1 (note the double-headed arrow), which sends any IWAlerts it generates to the former server.
- Both ASW Hubs relay their data to the ASW Console located on the Admin workstation (the dashed lines).

FIGURE 4-1: ASW in a seven-server environment

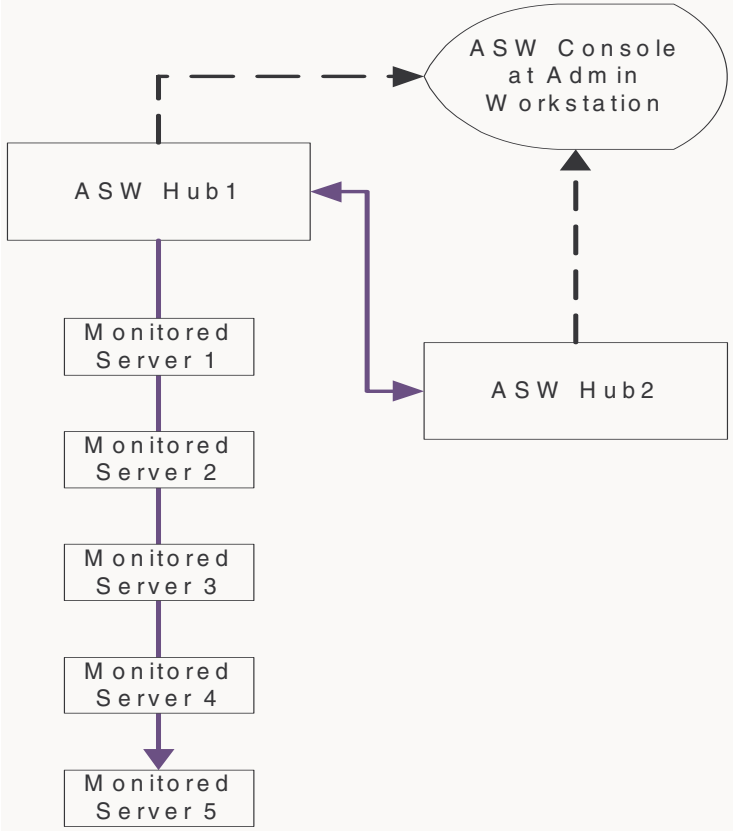

4.1.3.0 Scenario 2:

ASW deployment in an International Environment

Unlike Scenario 1, where all servers are in one geographical location, Scenario 2 presents the challenge of monitoring servers all over the globe.

Perhaps the foremost challenge is how to avoid the 'false alerts' caused by varying network response times. One of the three principles of ASW deployment mentioned above is especially applicable here: *ASW Hubs should monitor servers in their own geographical area*.

Note the following deployment features in Figure 4-2 on page 159:

- ASW Server is installed on the Primary Regional Hubs, ASWHub New York, ASWHub Singapore and ASWHub Capetown.

  Each of these Hubs monitors *all other servers in its respective region.*

- In addition, a second ASW Hub is installed *in each region*, ASW NYHub2, ASW SPHub2, ASW CTHub2.

  These secondary ASW Hubs monitor only the primary Hubs in their respective regions.

- The ASW Console, located on the central Admin workstation displays all six ASW Hubs, the servers they're monitoring, and all related messages/status info (see the dashed lines).

- Installing at least one instance of the ASW Console in each region facilitates your monitoring activities.

To avoid 'false alerts', determine the response times that can be expected in your environment, then make the necessary adjustments in the Network Time-out setting accessed under File > Preferences. (See Figure 4-7 on page 170)

FIGURE 4-2: ASW in an international environment

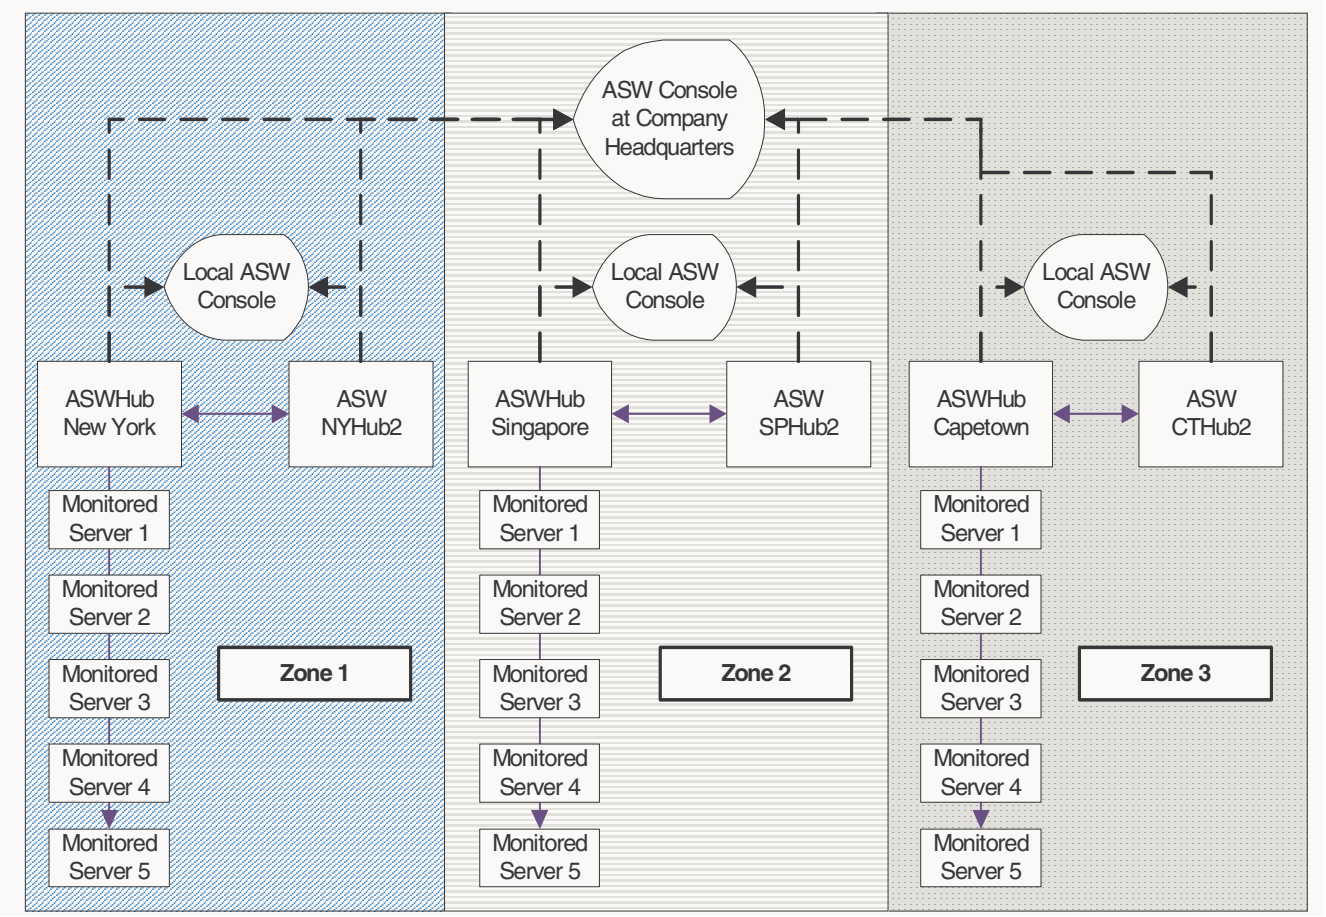

4.2.0.0 ASW BASICS

ASW is designed to carry out three basic functions:

- **check server connectivity and display current status**
- **notify of change of connectivity status**
- **display IWAlert messages by originating (monitored) server**

The first and last of these functions are quite straightforward, and need no elaboration here. (For details on the latter, see *"How can I view messages for a given server?" on page 169*).

The procedure by which ASW determines change of connectivity status is more complex, and bears further explanation.

4.2.1.0 Notify of change of status

ASW notification is based on a *change of status, rather than a condition of Not Responding*. What does this mean in practice?

ASW keeps track of two status variables: *Current* and *Previous*. When the Current status differs from the Previous status, configured notifications are sent (after working through any configured Retries). (See *"Action Profiles" on page 179*.)

For the sake of clarity, let's define these three critical terms:

- **Current**

Server status *as of the last monitoring cycle*.

- **Previous**

Server status *the last time notifications were sent*.

(The value of Current is assigned to Previous if *and only if* notifications have been sent.)

- **Retries**

Maintenance Release 27.33.

Retries only comes into play when:

- Previous state is Responding, and
- Current state is Not Responding

Each time the checked server is found to be Not Responding, Retries is decremented. When Retries equals zero *at the start of the monitoring cycle*, notifications are sent. (Table 4-1 on page 161 illustrates how Retries work. Note especially 1) what information is updated, and 2) how notification takes place when Retries reaches zero (0).

(To keep the following example short, Retries was set to one (1). Depending on the role of a given server in your environment, you may want to configure a higher number of Retries.).

Table 4-1. Retries parameter and notification

Stage of Monitoring Cycle	Previous	Current	State Changed?	Retries	Actions carried out
Start	Responding	Responding	No	1	■ No actions
Finish	Responding	Responding		1	
Start	Responding	Not Responding	Yes	1	■ Decrement Retries from 1 to 0
Finish	Responding	Not Responding		0	■ *Note: Current and Previous values remain unaffected. Previous state only updated after Notification takes place.*
Start	Responding	Not Responding	Yes	0	■ Send Notifications ■ Assign value of Current state to Previous (now both Not Responding) ■ Reset Retries
Finish	Not Responding	Not Responding		1	
Start	Not Responding	Responding	Yes	1	■ Notify immediately ■ Assign value of Current state to Previous (now both Responding) ■ Retries unaffected. It is bypassed when Not Responding changes to Responding
Finish	Responding	Responding		1	

4.3.0.0 ASW UNDER THE HOOD

The following flow chart allow you to predict what actions will be taken by ASW, depending on the server's Current and Previous states.

Remember ...

ASW keys off a Change of State, rather than a Not Responding condition.

If the Change of State is from Not Responding to Responding, configured notifications are sent immediately.

If the Change of State is from Responding to Not Responding, when notifications are sent depends on the value of the Retries parameter.

FIGURE 4-3: ASW program flow

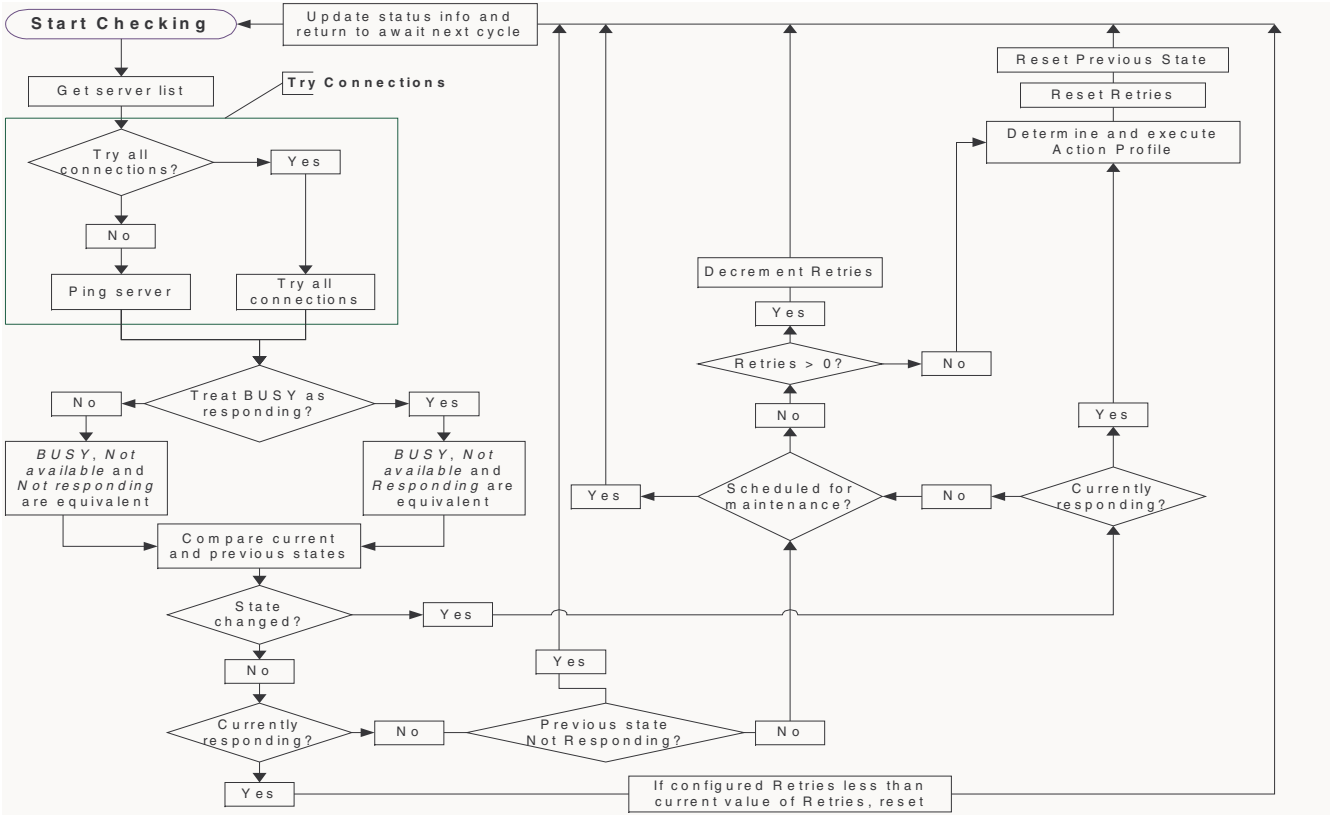

4.4.0.0 ASW CONSOLE

The ASW Console offers two views: Standard and Summary (see *Figure 4-4* and *Figure 4-5*).

FIGURE 4-4: ASW Console: Standard View

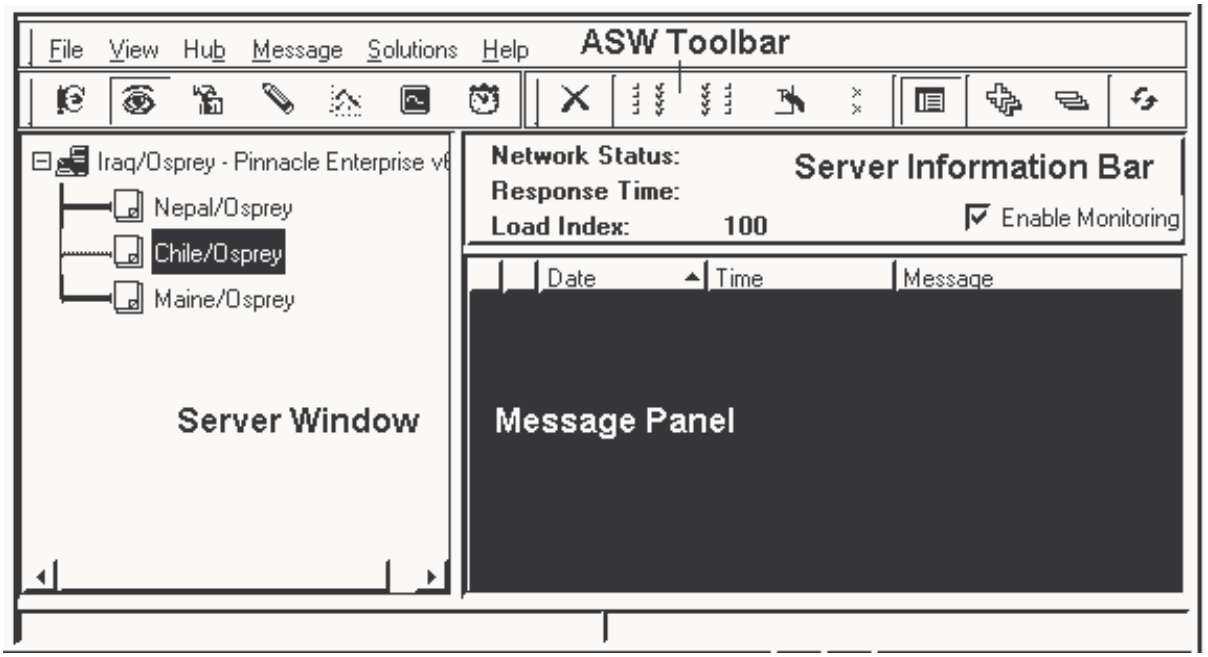

FIGURE 4-5: ASW Console: Summary View

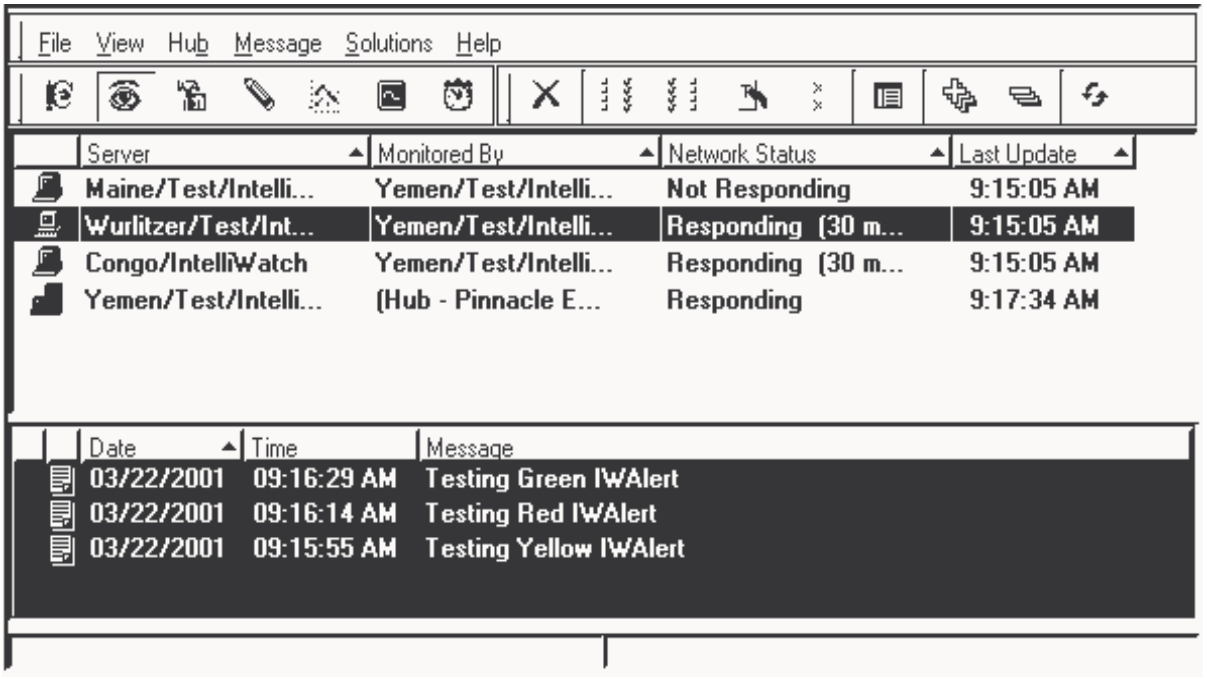

4.4.1.0 Standard vs Summary view

Principal differences of the two views are:

4.4.1.1 Standard

- Hubs/servers displayed in tree view
- Monitoring for individual servers can be turned off by unchecking Enabled box
- Includes Notes Load Index

4.4.1.2 Summary (default)

- Hubs/servers displayed in database view
- Sort by column
- Includes Last Update (time) field
- Large Icon display option
  - Go to **View > Large Icons** to toggle on and off.

4.4.2.0 Context-sensitive menus

Advanced ServerWatch offers context-sensitive pop-up menus, which are accessed by right-clicking at various locations on the user interface. The menu displayed depends on the position of the cursor.

Those menus that are the same for both ASW views are discussed first, followed by those available at the Standard view. Pop-up menus for the Summary view are discussed last.

4.4.2.1 Common pop-up menus

Only two pop-up menus are the same for both views. Both are accessed by right-clicking in the Message Panel (see *4.4.5.0*, below), with the cursor positioned over:

- one or more messages)

  - If one or more messages are selected—and you right-clicked on one of those messages to launch the pop-up menu—the top two menu items are displayed in the plural (Messages).
  - The functionality of the bottom two items is the same as when accessed via the drop-down menu (**Messages > etc.**), or a toolbar icon.

If one or more messages is selected, the Delete and Archive menu items affect all selected message—assuming you launched the pop-up menu by right-clicking on a selected item.

In all other cases, only the item clicked on when the menu is launched is affected.

- a blank part of the Message Panel

  - The functionality of both menu items is the same as when accessed via the drop-down menu (**Messages > etc.**), or, respectively, via the  and  toolbar icons.

4.4.2.2 Standard view

Two pop-up menus are available in the Standard view, both within the Server Window (see *4.4.4.0*, below).

The two menus are displayed with the cursor positioned, respectively, over:

- a Hub/monitored server

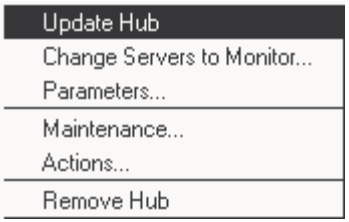

  – The functionality of all menu items is the same as when accessed via the drop-down menu (**Hub > etc.**), or the respective toolbar icon. See those sections for details.

Whether positioned over a Hub or a monitored server, the menu items affect the Hub.

- a blank area of the Server Window

  – The functionality of both menu items is the same as when accessed via the drop-down menu (**Hub > etc.**).

4.4.2.3 Summary view

The corresponding pop-up menus in the Summary view display the same items as the Standard view, in addition to the three discussed below.

The two menus are displayed with the cursor positioned, respectively, over:

- a Hub/monitored server

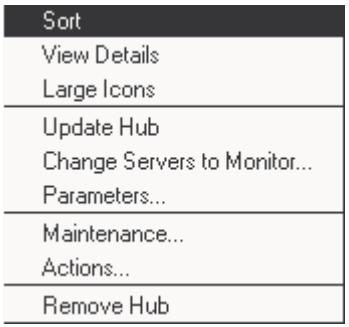

  – Sort

  The sort order depends on which column the cursor is over when you right-click.

  SERVER ICON: sorted in order of decreasing severity status (factors in presence and severity of any IWAlerts present).

  SERVER: sorted in ascending alphabetical order (disregards whether system is a Hub or a monitored server).

  MONITORED BY: sorted in ascending alphabetical order, listing Hubs first, then monitored servers.

  NETWORK STATUS: sorted in order of decreasing severity status.

  Servers not currently being monitored are listed *last*.

  LAST UPDATE: sorted in descending order, by date.

  Servers not currently being monitored are listed *first*.

*As a rule, the Last Update refers to the time/date recorded in the server document in **iwasw.nsf** (on the Hub), which represents when the server was last checked by the add-in task.*

If, however, a server document is modified (and saved) between monitoring cycles, the time/date stamp is updated. In that case, until the next monitoring cycle, the time/date displayed at the Pinnacle Console represents when those modifications were saved, and not the last time the server was checked by Advanced ServerWatch.

- View Details

  Click on this item to select/deselect the Standard view.

  The same effect is achieved either by clicking on the toolbar icon, or by going to **View > Summary/Standard View** via the drop-down menus.

- Large Icons

  Select to display larger Hub/monitored server icons.

  The same effect is achieved by going to **View > Large Icons** via the drop-down menus.

The functionality of all other menu items is the same as when accessed via the drop-down menu (**Hub > etc.**), or the relevant toolbar icon.

- over a blank area of the Server Window

For the functionality of the displayed items, see sections *4.4.2.2* and *4.4.2.3*, above.

4.4.3.0 Console Sections

The following discussion focuses on the Standard View of the ASW Console.

The Console is comprised of four sections:

- **Server Window**
- **Message Panel**
- **Server Information Bar**
- **ASW Toolbar**

4.4.4.0 Server Window

The Server Window displays three different icons, which in turn are color-coded at three levels:

- **Hub Icon**
- **Monitored Servers Icon**
- **Connection Symbol**

4.4.4.1 The Hub Icon

Only with the Hub server is the color of the icon affected by both message and connection status. In other words, for the Hub to be green, *both* of the following conditions must be met:

- *All* IWAlerts for *all* monitored servers are green (status messages).
- *All* servers being monitored are responding within the specified threshold

Table 4-2. Hub Server icons by color

Color of Hub Server Icon	What It tells you
Red 	An error condition exists *on one or more* of the servers being monitored by the Hub.
Yellow 	A warning condition exists *on one or more* of the servers being monitored by the Hub.
Green 	Only status messages exist *for all servers* being monitored by the Hub.
Gray 	No messages exist *for any of the servers* being monitored by the Hub, and *no* servers are being monitored.

4.4.4.2 The Monitored Servers Icon

Server icons display in the color of the *highest priority message* of the managed server. The highlighted server at Figure 4-6 on page 169, below, has sent an Error Message to the Console. Until the condition is corrected (and the message overwritten with an IWAlert to that effect), this server icon remains red.

Table 4-3. Monitored server icons by color

Color of Server Icon (not a Hub)	What It tells you
Red 	At least one error message exists *for this server.*
Yellow 	At least one warning message exists *for this server.*
Green 	Only status messages exist *for this server.*
Gray 	No messages exist *for this server.*

4.4.4.3 The Connection Symbol

Represents response time as of the last monitoring cycle.

Table 4-4. Connection symbols by color

Color of Connection Symbol	What It tells you
Red 	Server is not responding.
Yellow	Server is responding, but response time is slower than the designated threshold.

Table 4-4. Connection symbols by color (continued)

Color of Connection Symbol	What It tells you
Green	Server is responding within the specified threshold.
Gray	Server: is not being monitored; is being monitored but has not yet been checked; is down for maintenance.

You may see a red Server Icon, but a green Connection Symbol. This means that server-response time is under the threshold, but at least one IWAlert Error Message has been sent by that server to the Console, as illustrated at Figure 4-9 on page 173.

4.4.5.0 Message Panel

The Message Panel is located in the lower right hand corner of the ASW Console (see Figure 4-4, “ASW Console: Standard View,” on page 163, and Figure 4-6, “Detail of ASW Message Panel,” on page 169).

4.4.5.1 How can I view messages for a given server?

In the Server Window, select the managed server for which you want to see messages.

The Message Panel contains messages *for that server only.*

4.4.5.2 What do the message colors signify?

- **Green (1):** Status messages (for example, Responding)
- Yellow (2): Warning messages (for example, Server Sessions Dropped is [value])
- **Red (3):** Critical (error) messages (for example, Not Responding)

FIGURE 4-6: Detail of ASW Message Panel

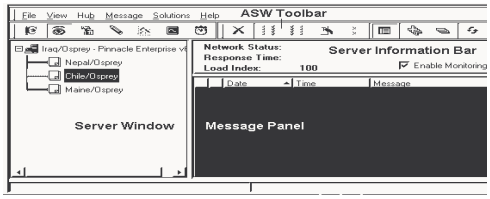

4.4.5.3 Why do certain messages disappear?

Two common reasons are:

- Message pairing: Messages reporting an error condition are overwritten when a message is received saying the condition no longer exists.

For message pairing to work properly, a Trigger cannot send IWAlerts to more than one Hub server.

- “Number of messages to return” parameter has been exceeded.

FIGURE 4-7: Preferences dialog

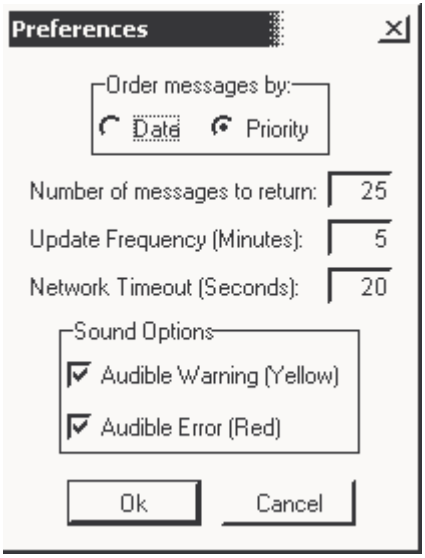

The ASW Console is set by default to display messages sorted by Date, and to show only the last 25 messages. To display messages by Priority, or to adjust the number of messages to return, go to **File > Preferences**, and make the desired configuration changes in the dialog pictured in *Figure 4-7*, above.

4.4.5.4 New in 27.36

Maintenance release 27.36 contains an enhancement that enables you to automate archiving and deletion of IWAlerts. See *4.9.0.0 on page 186*.

4.4.5.5 Audible warnings

Advanced ServerWatch gives you the option of selecting audible warnings for both yellow and red IWAlert messages (see *Figure 4-7*). The sound option requires that your machine be equipped with a sound card.

4.4.5.6 How can I tell which Monitor Trigger sent an IWAlert?

Simply double-click the message (or select the message and click the Display Trigger icon on the Toolbar).

4.4.6.0 Server Information Bar

The ASW Server Information Bar is located in the upper right-hand corner of the Console, and displays details about the connection status of individual managed servers. (See Figure 4-4, "ASW Console: Standard View," on page 163, and Figure 4-8, "ASW Server Information Bar," on page 171.)

To see information for a particular server, simply click its server icon in the Server Window. The fields are:

- **Network Status:** The server's current state (Responding, Not Responding. for example).
- **Response Time:** Value in milliseconds, as of the last monitoring cycle.
- **Load Index:** Based on Notes's Load Index, which rates Server Load on a scale of 1 to 100 (with 1 indicating *overloaded* and 100 *no load*).
- **Enable Monitoring:** If checked, indicates the server is currently selected for monitoring. To stop monitoring a server, select it in the Server Window, then uncheck this box.

  A changed Enabled status is not reflected in the Network Status field until the next time ASW reads the database

(at the next monitoring cycle). Therefore, for a short interval, you may notice an apparent conflict between these two fields.

FIGURE 4-8: ASW Server Information Bar

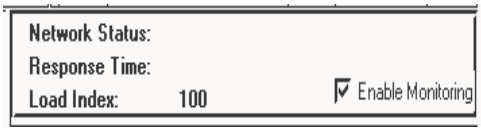

4.4.7.0 ASW Toolbar

The ASW Toolbar, located above the server window, simplifies a number of actions t the ASW Console. See the following table for details

Table 4-5. ASW Toolbar icons

Icon	Description
	This icon becomes active when a Hub Server is selected. Clicking on it *deletes the Hub Server*.
	*Selects all messages* currently displayed.
	*Deselects all messages* currently displayed.
	*Edit or view the trigger which sent an IWAlert.* (Before clicking on the icon, select the message. *This icon is enabled only if the message originated from a Management Agent trigger.*)

Table 4-5. ASW Toolbar icons

Icon	Description
	Deletes all selected messages.
	Switches from Standard to Summary View, and back.
	Updates messages from selected server.

When a Hub is selected, clicking the Refresh button refreshes all data; when a managed server is highlighted, however, only IWAlert messages are updated.

4.5.0.0 ASW ACTIVITY LOGGING

Starting with IntelliWatch, the response time of servers being monitored by Advanced ServerWatch can be logged to **iwstats.nsf**, for later reporting using IntelliWatch Analyzer.

4.5.1.0 Enabling activity logging

Activity logging is turned on/off at the Hub level, *not* at the monitored-server level.

TO ENABLE ACTIVITY LOGGING:

1 In the left-hand pane of the Advanced ServerWatch Solution, select either the Hub for which Activity Logging is to be turned on/off, or any one of it's monitored servers.

2 Go to **Hub > Monitoring > Parameters** via the drop-down menus.

3 To turn on Activity Logging, select the *Enable Activity Logging* checkbox; to turn off Activity Logging, deselect the checkbox.

Response times will now be logged to the database.

Statistics are logged with the format:

ResponseTime.from.<hubserver>.to.<pinged server> = <time in milliseconds>

Response-time statistics can be reported on using Analyzer.

4.6.0.0 WORKING WITH ASW HUBS

4.6.1.0 Setting up ASW Hubs

When setting up ASW Hubs, you are not creating the add-in task itself, but establishing a connection between it and the ASW interface at the Pinnacle Console.

4.6.1.1 What occurs when you "Add a new hub" at the Console?

- **Console attempts connection to ASW task on Hub to be added**

  – If the attempt is successful (meaning the server is available and the ASW add-in task is loaded), the Hub appears in the server window of the ASW Solution at the Pinnacle Console.

- **Console displays monitored servers already configured on Hub**

  – Even though your workstation may be connecting to a given ASW Hub for the first time, the Hub may already have been configured to monitor servers (via another connection, or directly in the Advanced ServerWatch database **iwasw.nsf**). Those servers are displayed without your having to add them.

4.6.1.2 Security Mechanism

For the security mechanism governing access to the Console, see *1.4.0 on page 24*.

4.6.1.3 Connecting to ASW Hubs at the Pinnacle Console

To manage your ASW Hubs, bring up the Pinnacle Console, then select the ASW Solution, using the following steps.

The following procedure assumes you are accessing the Pinnacle Console via Internet Explorer. If you are using the stand-alone client, skip Steps 1 and 2.

TO CONNECT TO ASW HUBS:

1 Make sure your Primary Server is up, and the HTTP task is loaded.

2 Bring up Internet Explorer, and enter a URL for your Primary Server based on the following example:

```
http://[DomainNameOfPrimaryServer]
/[IntelliWatchDataDir]/console.ns
f
```

Depending on the configuration of your domain (and from where the Primary Server is being accessed), you may need to use the complete Internet domain name.

For example: If the Host Name of your Primary Server is MyServer/MyDomain, from within the domain, "MyServer" should suffice. Under some conditions, however, you may need to enter "MyServer.MyDomain.com" Or, to return to the example under Step 3, above:

http://MyServer.MyDomain.com/[IntelliWatchDataDir]/console.nsf

3 Make sure the **iwasw** task is running on any Hubs you want the console to display.

4 Go to **Solutions > Adv. ServerWatch** via the drop-down menus; alternatively, click on the  toolbar icon.

4.6.1.4 Adding a new ASW Hub

The following procedure assumes you have already followed the steps at *"Connecting to ASW Hubs at the Pinnacle Console" on page 172*.

TO ADD AN ASW HUB:

1 From the Hub menu, select "Add a new hub".

2 In the dialog brought up by the previous step, click on the combo box arrow and select from the list displayed, or start typing.

- As you type, the name in the text box is the first in the server list that corresponds to all the letters typed so far.
- If the Domino server name differs from the host name (or Title), fill in the server name, then click the Advanced button and put the host name in the Title field.

3 Click OK to add the Hub.

4 Repeat the above steps for each Hub.

If you try to connect to the ASW Hub task on a machine where that task is not running, the error message in the following figure is displayed.

You may continue with Hub setup, or abort the operation. If you choose to continue, the Hub is displayed with a question mark over the icon.

If you (re)start the iwasw task on that server, the Hub icon returns to normal the next time the server is checked by the ASW console.

FIGURE 4-9: Error message when iwasw not running on Hub

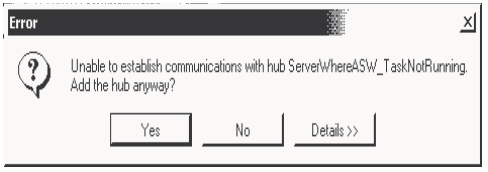

4.6.1.5 Deleting ASW Hubs

The following procedure assumes you have already followed the steps at *"Connecting to ASW Hubs at the Pinnacle Console" on page 172*.

TO DELETE AN ASW HUB:

1 Select the Hub to be deleted in the Server Window.

This activates the  toolbar icon.

2 Click on the delete icon.

This brings up a "Remove Hub" dialog.

3 To delete the Hub, hit Enter, or click on the Yes button.

Deleting an ASW Hub at the Pinnacle Console has no effect on the ASW Add-in task running on the server in question. That instance of ASW Server continues to monitor the servers assigned to it, and no information is lost. The Hub just won't be visible at the Pinnacle Console.

Deleting server documents in an ASW Hub's iwasw.nsf does cause monitoring to cease for the relevant managed servers.

4.6.2.0 Enable Monitoring

The Advanced ServerWatch task monitors every server for which there is a corresponding Server Document in the Server List View of the **iwasw.nsf**.

If there is no corresponding server document in that Hub's ASW database, a server will not be monitored by that instance of Advanced ServerWatch.

Server documents can be created in three ways:

- check Enable Monitoring on the Server Information Bar (see *4.6.2.1*, below)
  - Applies only if the server has sent IWAlert messages to the Hub in question.
- select server via the Pinnacle Console (see *4.6.2.2*, below)
  - Method for servers that have not sent IWAlert messages to the Hub in question (although it can be used for all servers).
- directly in the database (see *10.3.2.1*)

4.6.2.1 Servers that have sent IWAlerts

Even when a server has no corresponding server document in a Hub's ASW database, an icon for that server is displayed in the Server Window under that Hub if it has sent IWAlert messages to the Hub.

(The IWAlert messages are displayed in the Message Panel to the right of the Server Window.)

Since no server document for that system exists in that Hub's **iwasw.nsf**, the Enable Monitoring checkbox is not selected.

To create a server document for this system in the ASW database on the Hub, simply check Enable Monitoring (on the Server Information Bar).

(The document is created immediately in **iwasw.nsf**.)

4.6.2.2 Servers that have NOT sent IWAlerts

Servers that have no corresponding server document in a Hub's ASW database, and that have NOT sent IWAlert messages to the ASW Hub, are not displayed as icons at the

Pinnacle Console. Therefore, the server cannot be selected, and a server document created, using the method discussed at *4.6.2.1*, above.

In those cases, a server document must be created either directly in the database (see *10.3.2.1*), or by the following procedure.

(This procedure assumes you have already followed the steps at *"Connecting to ASW Hubs at the Pinnacle Console" on page 172*.)

TO ADD MANAGED SERVERS:

1 In the left-hand pane of the ASW Console, highlight the Hub for which you want to add servers.

2 Go to **Hub > Monitoring > Change Servers to Monitor...** via the drop-down menus.

This brings up the dialog "Select Servers to be Monitored by [servername]".

FIGURE 4-10: Managed Server Selection Dialog

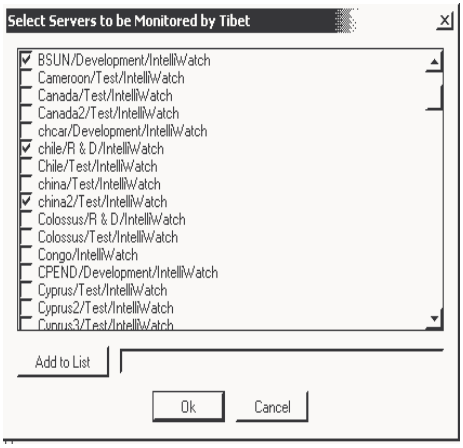

This dialog differs by design from other dialogs with a similar purpose (see, for example, Figure 4-13 on page 178). Functionality is not affected.

3 Select the servers you want the ASW Hub to monitor.

- If the server you want to monitor is in your environment, but does not show up in the list, type the name in the text box provided, then click the Add to List button.

4 Click OK.

This process creates a server document for each newly selected managed server in the iwasw.nsf database on the relevant ASW Hub.

5 Follow the same steps for each of your ASW Hubs.

4.6.2.3 Deleting Managed Servers for an ASW Hub

The following procedure assumes you have already followed the steps at *"Connecting to ASW Hubs at the Pinnacle Console" on page 172*.

Managed servers for which the ASW database contains messages cannot be deleted from the display, unless you first delete those messages from the database.

TO DELETE MANAGED SERVERS:

1 Click on the relevant ASW Hub (in the left-hand pane of the ASW Console).

See note under *"Enable Monitoring" on page 174*.

2 Go to **Hub > Monitoring > Change Servers to Monitor** via the drop-down menus.

This brings up a dialog that allows you to add or remove servers associated with the selected Hub.

3 Uncheck the server(s) you want to delete from the Managed Server List for this Hub.

4 Click OK to confirm the deletion.

4.7.0.0 MAINTENANCE AND ACTION PROFILES

ASW provides two types of server profiles:

- **Maintenance Profiles** allow you to 'inform' ASW Server of periods of scheduled maintenance or downtime.
- **Action Profiles** allow you to specify notification methods to take when the connection to a managed server is lost.

4.7.1.0 Maintenance Profiles

Servers down for maintenance are unable to respond to ASW connection probes. Maintenance Profiles ensure that these servers do not generate unwanted and unwarranted notifications.

Maintenance Profiles inform the ASW server task that specified servers are unavailable. When a server is deemed Not Responding, ASW checks to see if it is covered by a currently active Maintenance Profile. If it is, ASW converts the *Not Responding* status to *Scheduled for Maintenance*, and no actions are taken. (See *"ASW program flow" on page 162*

A default profile is created during the Setup with a wildcard (*) in the Server Groups field. Every ASW Hub you create starts with this default profile, and all (enabled) managed servers under that Hub are governed by it.

Additional profiles are created *for individual ASW Hubs*, so be sure to select the relevant Hub *before creating or editing a profile*.

Maintenance Profile fields are self-explanatory, with two exceptions: Interval and Subinterval.

4.7.1.1 Interval and Subinterval fields

Maintenance Profiles are listed with both an Interval and a Subinterval field (see the third and fourth column headings at Figure 4-11, "Maintenance Profiles dialog," on page 177).

The **Interval** field contains one of three available options:

- Daily
- Weekly
- Monthly
- Day of Month

The Interval is selected via a radio button on the *"Chapter 4, Create new Maintenance Profile dialog"* (see Figure 4-12 on page 178).

Subinterval does not appear as such on that dialog, and is populated differently, depending on the *Interval* selected:

- Daily: The Subinterval field is blank, and no selection is required (nor indeed possible) on the *"Chapter 4, Create new Maintenance Profile dialog"*.
- Weekly: Selecting this option causes check boxes to appear for the seven days of the week. The Subinterval column of the *"Chapter 4, Maintenance Profiles dialog"* is populated with the abbreviated form of the days selected (see Figure 4-11 on page 177).
- Monthly: Selecting this option on the *"Chapter 4, Create new Maintenance Profile dialog"* causes a text box to appear, in which you enter the day of the month *as an integer* (1-32).

  This integer subsequently shows up in the Subinterval column of the *"Chapter 4, Maintenance Profiles dialog"* (see the Subinterval column of the highlighted profile at Figure 4-14 on page 179).

  The integer 32 activates the profile on the last day of the month, regardless of its length.
- Day of Month: Selecting this option on the *"Chapter 4, Create new Maintenance Profile dialog"* causes two sets of checkboxes to appear:

  – use column one to select the week(s) of the month

  – use column two to select the day(s) of the selected weeks for which to activate the profile

  The profile at Figure 4-12 on page 178schedules maintenance for the first and third Thursdays of every month, from 8:30-10:30pm, for example.

4.7.1.2 Creating Maintenance Profiles

Although the second text box on the Create New Maintenance Profile dialog is entitled *Server Group*, individual servers may be entered. Nevertheless, we recommend using group names.

Especially if your server groups have short, descriptive names, a list of groups likely means more to you—and be more compact—than a list of individual servers.

The following procedure assumes you have already followed the steps at *"Connecting to ASW Hubs at the Pinnacle Console" on page 172*.

TO CREATE MAINTENANCE PROFILES:

1 Select the Hub server for which the profile is being created.

2 Select **Hub > Profiles > Maintenance** from the menu bar. This displays a dialog listing Maintenance Profiles. At first, you'll only see the default profile created by the Setup (see *Figure 4-11*, below).

FIGURE 4-11: Maintenance Profiles dialog

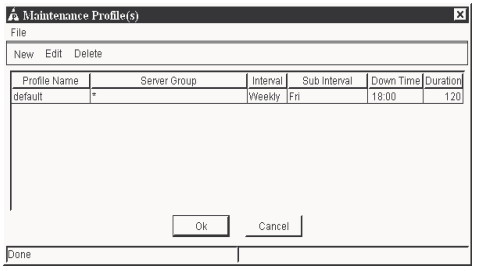

3 Click on the New button, or go to **File > New** via the drop-down menus. The *"Chapter 4, Create new Maintenance Profile dialog"*, below, is displayed).

FIGURE 4-12: Create new Maintenance Profile dialog

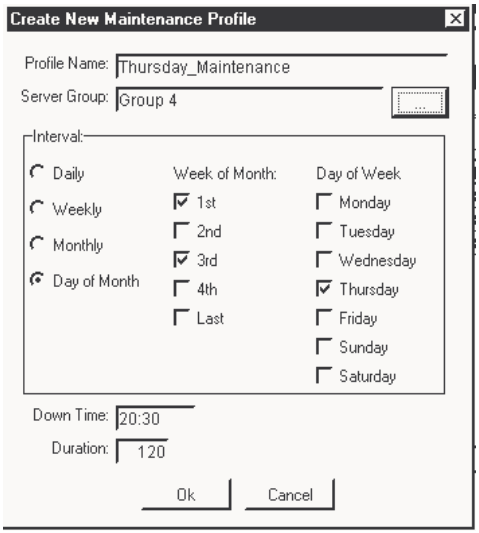

4 When specifying the Profile Name, we suggest you use a name clearly descriptive of the Profile's purpose, such as *NTCentral/Group XYZ Weekly Maintenance*.

5 The Server Groups field can be filled in by hand, or by selecting the button to the right of that field (which brings up the dialog displayed in *Figure 4-13*, below).

- Highlighting one or more Choices of servers/groups from the left-hand list box enables the Add button, allowing you to move them to the Selected list.
- If you want to remove a server or group after selecting it, simply highlight it (which enables the Remove button), and click Remove.

FIGURE 4-13: Select servers dialog

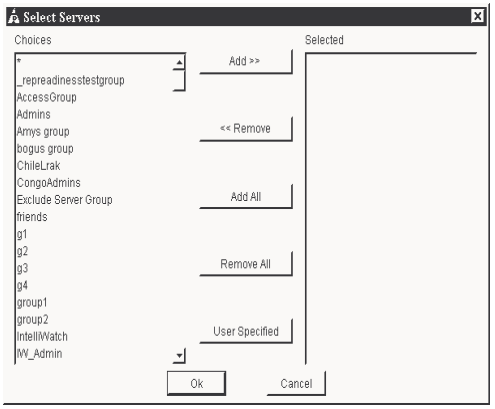

6 The Down Time field (see Figure 4-12 on page 178) accepts both 24-hour and 12-hour formats. If A.M. or P.M. is not specified, 24 hour time is assumed.

7 Click OK to save the Profile.

4.7.1.3 Editing Maintenance Profiles

The following procedure assumes you have already followed the steps at *"Connecting to ASW Hubs at the Pinnacle Console" on page 172*.

TO EDIT MAINTENANCE PROFILES:

1 Select **Hub > Profiles > Maintenance** from the menu bar.

This displays a dialog listing Maintenance Profiles. Highlight then double-click the profile you want to edit; alternatively,

2 Highlight the profile and click the Edit button on the toolbar.

See *Figure 4-14*, below.

3 Make any required changes, then click OK to save the Profile.

FIGURE 4-14: Selecting Maintenance Profile for editing

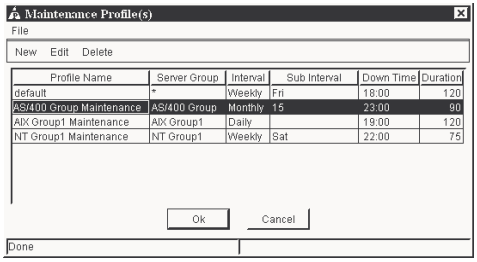

4.7.2.0 Action Profiles

Action profiles specify what to do when a managed server appears unreachable.

A default profile is created during the Setup with a wildcard (*) in the Server Groups field. Every ASW Hub you create starts with this default profile, and all (enabled) managed servers under that Hub are governed by it.

Profiles are created for each Hub Server independently, so be sure to select the target Hub Server before creating or editing a profile.

Before proceeding with the creation of an Action profile, please note the following regarding the Notification options available by platform (see *Figure 4-16*, below):

- **E-mail:** Works on all platforms.
- **Start Program:** Works on all platforms (for usage details, see *4.7.4.0 on page 184*).
- **Page:** Works on all platforms. *However*, the Paging Server itself must be installed on an NT Workstation or Server.
- **NT Event Log:** Works only on NT.
- **SNMP Trap:** Works on all platforms.

4.7.2.1 Enhanced Action Profiles (27.33 and above)

Prior to maintenance release 27.33, only one Action Profile could govern ASW's response to a loss of connectivity between the Hub and any given (monitored) Domino server. This was true whether the server was listed explicitly in the profile, or was a member of a listed group.

You could always CREATE numerous Action Profiles that applied to a given server/group, BUT, in the event that ASW deemed a server to be NOT RESPONDING, only ONE of those profiles would be executed.

Since the mechanism that determined which Action Profile would be executed has been changed, details as to how it worked previously are not relevant here.

Details...

- **Profile name need not be unique**

Previously, it was not possible to save an Action Profile with the same name as an existing profile--even if the new profile was configured differently.

Starting with maintenance release 27.33, ASW allows you to create more than one Action Profile with the same name--provided the number of Retries is different. (If the number of Retries is the same, but the

Profile is otherwise configured differently, you will not be allowed to save it.)

- **Escalating Action Profiles**

The principal advantage of the latest enhancements to Action Profiles is that they allow you to take a series of escalating steps, the longer a server remains unavailable.

Imagine you have a server that goes down after-hours. As long as the server is back up in, say, less than 10 minutes, you only want to notify the 24x7 Help Desk, and log an NT Event. Assuming your ASW Hub has a Monitoring Frequency of 5 minutes, you could create an Action Profile with Retries set to zero (0), and with the Help Desk Group as the only recipient of any notifications. Let's call this profile XYZ_Alert.

Now imagine that, if this server remains unavailable for an additional 10 or 15 minutes, you want to page the Admin On-Call (as well as carry out certain other actions). Simply create a new Action Profile, set Retries to 2, and configure it to carry out the desired actions. Let's call this profile XYZ_Caution.

Lastly, if the server remains unavailable for more than 30 minutes, you want to page the Head Admin. You might create a third Action Profile, set Retries to 5, and configure the profile to carry out the desired actions. We'll call this last profile XYZ_Critical.

4.7.2.2 What happens when the monitored server is again RESPONDING?

In previous product releases, when ASW detected that a monitored server was again RESPONDING, the relevant Action Profile sent notifications to all e-mail addressees that had been notified of a loss of server availability (as well as carrying out any other configured actions).

If the list of e-mail addressees contained a user both explicitly and as part of a group, that person would receive two notifications.

What happens with the enhanced Action Profiles, since a given addressee (or other action) may be included in more than one executing profile? Enhanced Action Profiles keep track of all addressees who receive notifications regarding a particular server, as well as of any other configured actions--in all applicable profiles. When the relevant server is again RESPONDING, each notification/other action will be executed only ONCE--provided an explicitly named addressee of a notification action is not part of a group to which notifications are also sent.

4.7.2.3 A practical example

Let's say the three imaginary Action Profiles outlined above were configured progressively to take the following actions:

- XYZ_Alert
  - E-mail the Help Desk group
  - Log an NT Event
- XYZ_Caution
  - Page the On-Call Admin
  - E-mail the Head Admin
  - E-mail the Help Desk group
  - Log an NT Event
  - Run Batch file 1
- XYZ_Critical
  - Page the Head Admin
  - E-mail the On-Call Admin
  - E-mail John Smith

    (Help Desk manager and member of Help Desk Group)
  - Log an NT Event
  - Run Batch file 1

When a server covered by these profiles comes back up, the net result will be the same as the current behavior, that is to say that the following actions will be taken ONCE:

- E-mail
  - E-mail the Help Desk group

    E-mailed only ONCE, even though two Profiles contain this action (P1 & P2).
  - E-mail the On-Call Admin
  - E-mail the Head Admin
  - E-mail Joe Smith

    He will get 2 messages, one as a member of the Help Desk Group, and one as an explicitly named user.
- Pages
  - Page the On-Call Admin
  - Page the Head Admin
- Other actions
  - Log an NT Event

    Executed only ONCE, even though all Profiles contain this action.
  - Run Batch file 1

    Executed only ONCE, even though two Profiles contain this action (P2 & P3).

Clearly, more profiles can be created to further enhance the escalation process, but these three profiles illustrate what's possible.

4.7.2.4 Creating Action Profiles

The following procedure assumes you have already followed the steps at *"Connecting to ASW Hubs at the Pinnacle Console" on page 172*.

TO CREATE ACTION PROFILES:

1. Select the Hub server for which the profile is being created.
2. Select **Hub > Profiles > Actions** from the menu bar.

This displays a dialog listing Action Profiles. At first, you'll only see the default profile created by the Setup (see *Figure 4-15*, below).

3. At the *"Chapter 4, Action Profiles dialog"*, either click on the New button, or select **Options > New** via the drop-down menu.

The dialog at *Figure 4-16*, below, is displayed.

FIGURE 4-15: Action Profiles dialog

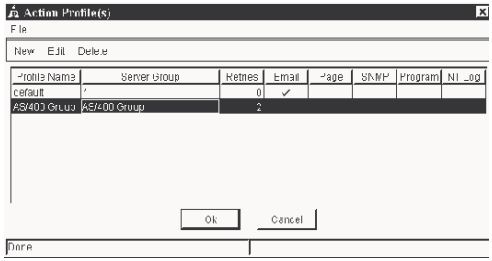

FIGURE 4-16: Create new Action Profile dialog

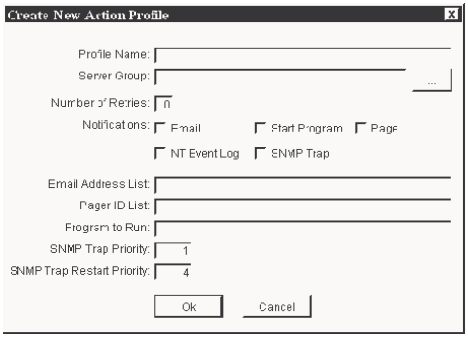

4 Using the check boxes, select those actions you want ASW to take when the Hub Server loses connectivity with a managed server.

5 When specifying the Profile Name, we suggest you use a name clearly descriptive of the Profile's purpose, such as *NTGroup1Notification*.

6 Server Groups can be filled in by hand, or by selecting the button to the right of that field.

This brings up the dialog at Figure 4-13 on page 178.

- Highlighting one or more Choices of servers/groups from the left-hand list box enables the Add button, allowing you to move them to the Selected list.
- If you want to remove a server or group after selecting it, simply highlight it (which enables the Remove button), and click Remove.

7 Select OK to save the profile.

4.7.2.5 Editing Action Profiles

The following procedure assumes you have already followed the steps at *"Connecting to ASW Hubs at the Pinnacle Console" on page 172*.

TO EDIT ACTION PROFILES:

1 Select **Hub > Profiles > Action** from the menu bar. This displays a dialog listing Action Profiles.

2 Highlight then double-click the profile you want to edit; alternatively, highlight the profile and click the Edit button on the toolbar.

See *Figure 4-17*, below.

3 Make any required changes, then click OK to save the Profile.

FIGURE 4-17: Selecting Action Profile for editing

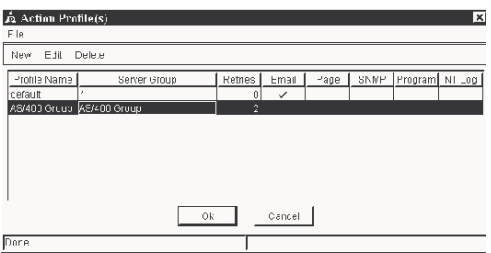

4.7.2.6 Treat BUSY as Responding Option

When ASW pings a server, it receives one of four responses:

- RESPONDING
- NOT_RESPONDING
- BUSY
- **RESTRICTED**.

The **Treat BUSY as responding** option concerns the last two of these: BUSY and RESTRICTED.

If this option is set to **True**, BUSY and RESTRICTED *are treated the same as RESPONDING,* that is, *no actions are taken*.

If the option is set to **False**, BUSY and RESTRICTED *are treated the same as NOT_RESPONDING,* that is, *actions are taken*.

This option can also be configured via a Notes client. For details, see *"Treat BUSY as responding" on page 337* of *"Chapter 10, Configuring IntelliWatch via Notes"*.

4.7.3.0 May **iwasw.nsf** replicate?

As installed by the Setup, the ASW database may not--and indeed can not--replicate.

Why is this?

Each copy of **iwasw.nsf** is created with a separate Replica ID, making replication impossible.

Why was this done?

Because replicating more than a specific subset of this database's documents can lead to issues on your ASW Hubs.

4.7.3.1 What does **iwasw.nsf** contain?

The ASW database contains the following views:

- Messages
  - These IWAlerts are sent by individual servers to the Hub, for display at the Pinnacle Console.
- Monitoring Parameters
  - These parameters govern:
    - monitoring frequency
    - response threshold
    - enabling activity logging
    - activity logging frequency (if enabled)
- Maintenance Profiles (see *4.7.1.0*)
- Action Profiles (see *4.7.2.0*)
- Server List

Of the above, only Maintenance and Action Profiles may (but need not) replicate.

Replicating anything other than Maintenance and Action Profiles is not supported.

4.7.3.2 What if I want to replicate Profiles?

If you want your Maintenance and Action Profiles to replicate, you must carry out the following procedure.

For purposes of illustration, assume you have two ASW Hubs, called SourceHub and Target Hub:

TO REPLICATE PROFILES:

1. Delete the **iwasw.nsf** database on TargetHub.
2. Create a replica of the SourceHub copy of **iwasw.nsf** on TargetHub, using your Notes client.
3. Limit replication between the databases to Maintenance and Action Profiles, using the following formula.

```
SELECT((Form="Notification Profile")||
  (Form="Maintenance Profile"))
```

4.7.4.0 Environment variables

Previous versions of Advanced ServerWatch allowed a descriptive string about a server's state to be passed to a batch/script file as the first parameter.

Example 1:

Advanced ServerWatch is configured to run the file **action.bat** when it sends notifications.

Action.bat has the following contents:

```
echo %1 > c:\out.txt

exit
```

If Advanced ServerWatch detected that SERVERXYZ was not responding, it would pass c:\out.txt the string:

"SERVERXYZ was not responding"

Similarly, once the server again became available (and notifications were sent), this same batch file would pass the string:

"SERVERXYZ is responding".

In IntelliWatch, Advanced ServerWatch also sets two named environment variables:

- NOTES_STATUS
- NOTES_SERVER

NOTES_STATUS will be one of the following:

- RESPONDING
- THRESHOLD_EXCEEDED
- NOT_RESPONDING
- SCHEDULED
- BUSY
- RESTRICTED

NOTES_SERVER will be the name of the server that generated the message.

Example 2:

If the file **action.bat** had the following contents:

```
@ECHO OFF

if %NOTES_STATUS% ==
RESPONDING GOTO ELSE

REM [Do something here to inform
you that the server stops
responding]

echo %NOTES_SERVER% is not
responding > c:\out.txt

GOTO END

:ELSE

REM [Do something here when the
server becomes available again]

echo %NOTES_SERVER% is
responding > c:\out.txt

:END
```

The batch file in this example is only using the RESPONDING and NOT_RESPONDING variable options, but the same technique can be used to incorporate all possible values of NOTES_STATUS.

4.8.0.0 MONITOR MESSAGES TO ASW

4.8.1.0 "Last Evaluation" messages from Triggers to ASW

Starting in maintenance release 27.36, each time a Trigger evaluates, a special "I am alive" message creates (or updates) a document in **iwasw.nsf** which contains the server name and the date.

A throttle on this function ensures that a message will be sent at most once a minute (not user-configurable).

4.8.2.0 Changes to ASW

ASW will periodically check these messages (after each ASW-monitoring cycle completes) to make sure that they are older that a user-configured time span (default is 15 minutes). To access this setting, go to **Hub > Monitoring > Parameters…**, as shown in Figure4-18, below.

FIGURE 4-18: Accessing new ASW settings

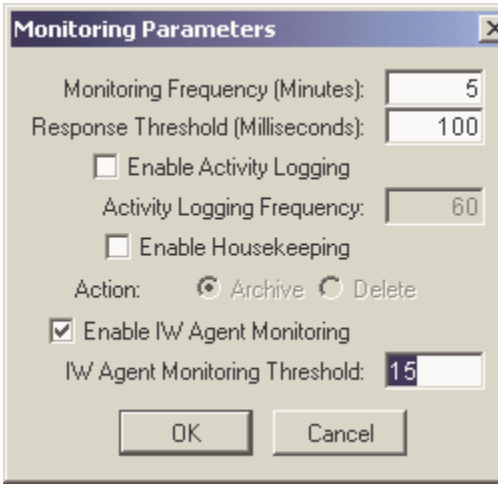

Deselecting the check box allows you to turn off this option.

When the user-specified time limit is exceeded, a red IWAlert message is created, and displayed at the ASW Solution of the PInnacle Console.

As with other IWAlerts sent by Triggers, once document history is again within the user-specified time limit, the red message will be updated. To see message history, right-click on the relevant IWAlert, as in Figure 4-19.

FIGURE 4-19: Last Evaluation History

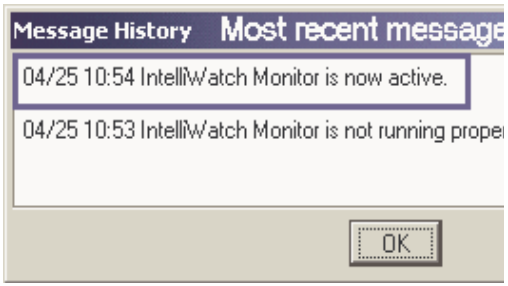

4.8.2.1 With pre-27.36 ASW Hubs

If IWAGENT is configured to send Last Evaluation messages to a pre 6.00.27.36 ASW Hub, the message will not be sent.

Instead, a message like the following will be issued to the iwagent.log file:

- "Unable to create Last Evaluation message because the Hub server is an older version (pre 6.00.27.36).").

4.9.0.0 AUTO-CLEANUP OF IWALERTS

4.9.1.0 Purpose of feature

To prevent the unnecessary buildup of IWAlert messages.

4.9.1.1 How to turn it on

Select the "Enable Housecleaning" checkbox at Figure 4-18.

Once the checkbox has been selected, the radio buttons (grayed out in the illustration) are enabled, allowing you to choose between archiving and deletion of the affected IWAlerts.

4.9.1.2 How it works

On startup, IWAGENT will send a special "Up" message to the hub server (if for some reason, one already exists, another one will be created). ASW will periodically process and delete these messages, to keep them from building up in significant numbers.

Once a minute, after each monitoring cycle, ASW checks for the special "Up" messages, then processes them as per user configuration (i.e. they are either archived, deleted, or ignored (if Enable Housecleaning is disabled--see Figure 4-18).

This setting is at the Hub level and not at the individual server level—all servers under a given Hub are treated alike.

Upon completion of processing the "Up" message will be deleted.

If no Up message is received, the IWAlert messages will never be cleaned up (to maintain existing behavior for previous versions of IW that may be running on spoke servers).

4.9.1.3 With pre-27.36 ASW Hubs

If IWAGENT is configured to send IWAlerts to a pre 6.00.27.36 ASW Hub, the message will not be sent. Instead, a message like the following will be issued to the iwagent.log file:

- "Unable to create Up message because the Hub server is an older version (pre 6.00.27.36).").

4.10.0.0 FAQs

Q: Why do no messages show up at the ASW Console for a particular server?

Example: Server 123 is selected under Hub ABC. The server is responding, but no messages are showing up at the Console. I have confirmed that the server is sending messages. Why can't I see them?

A: In all probability, the Hub Server parameter on Server 123 does not read Hub ABC. Please remember that an ASW Hub can monitor a server, even if that server is not reporting to the Hub in question. *However, the Hub does not receive messages from that server.*

To change the Hub Server setting, use the *"Chapter 11, Parameter Configuration Utility"* described on *page 350*.

Q: Why do the servers under an ASW Hub disappear?

Example: Hub DEF was monitoring several servers. All of a sudden, the servers disappeared, and the Hub icon is now overlaid by a ❓.

What happened?

A: Check to see that the **iwasw** task is still running on Hub DEF. Most likely, it's not.

Remember ...

*The Console gets the list of servers being monitored from the **iwasw** task on the Hub. If the task isn't running, the Console can't get the list, and displays the icon you saw.*

Q: Why does a server icon display a color for which there does not appear to be a matching colored message?

Example: The icon for a managed server is yellow, but when you drill down you find no yellow messages for that server. Where can that message be found?

A: The ASW Console is set by default to display messages sorted by date, *and* to show only the last 25 messages. If the yellow message concerned is not among the most recent 25 messages, you would not see it, *unless* you **a)** changed the sort method to priority rather than date, and/or **b)** increased the number of messages the ASW Console displays

Objective: To allow the user to select a maintenance profile that designates a specific time period for a given day of the week for a given week of the month. Ex: 10 – 12 PM on the third Thursday of every month.

Chapter

Performance Manager

5

Pinnacle Performance Manager was created to:

- allow customization of native Notes statistics
- supply statistic types not available in Notes.

PM statistics can be collected and used as data for report generation (by Analyzer). They can also form the basis of a condition to be monitored by IntelliWatch Triggers.

Chapter Contents

5.1.0.0 OVERVIEW

5.1.1.0 What makes PM statistics so useful?

Lotus Notes by itself keeps track of a number of useful statistics. Examples include the number of TCP/IP bytes received, the number of waiting mail messages, the number of users, and total mail routed. Notes' native statistics mechanism has two key drawbacks, however:

- statistics cannot be manipulated
- statistics reflect the count *since the last time the server was restarted*

Use PM to:

- customize Notes native statistics—for example by resetting them daily—without having to restart the Domino server
- create statistics of types not supplied by Notes
- compare, average and sum statistics
- generate statistics reports using IntelliWatch Analyzer
- trigger on user-specified statistic thresholds, using IntelliWatch Monitor

5.1.2.0 Enhancing Notes statistics

An excellent illustration of the power of PM statistics is the Delta type, which allows stats to be reset at a user-specified interval.

The fact that Notes native statistics reset only when the server is recycled complicates the interpretation of their significance. Since servers are unlikely to be recycled more often than once a week (and may have a much longer maintenance cycle, depending on the operating system), cumulative statistics become less rather than more meaningful as time goes on.

In addition, servers may need to be brought down (or they may crash) mid-cycle, making the comparison of statistics from one server to another difficult, even if one knows the 'availability history' of both systems.

All of the above requires Admins to 'filter' statistics by considering questions such as:

- What is the server's maintenance cycle?
- How many days has it been since maintenance was last performed?
- Has the server gone down in the meantime, and, if so, when (and for how long)?

Consider the following real-world example of Delta statistics in action.

Example 1: Delta statistic

On certain servers, the number of TCP/IP bytes received is abnormally high, especially during certain hours of the day. You have taken steps to bring it down, but want to see if these measures have been successful.

Using PM's *Delta* type, create a statistic that records TCP/IP bytes received, and resets that value hourly. A report based on this statistic will enable you to track the hourly fluctuations in this value, and make the necessary adjustments in your infrastructure.

5.1.3.0 PM Statistic architecture

To effectively create Performance Manager statistics, and report on them using Analyzer, it's essential to understand some basic facts about IntelliWatch statistical architecture:

5.1.3.1 PM's built-in types

Each PM statistic type returns very specific data.

In most cases, PM appends one or more elements to the Statistic Name you create. The *Mail Outgoing Attachment Type*, for example, returns the following:

[Statistic Name].[extension].Count
[Statistic Name].[extension].Size

Some PM statistics—the Delta type is a case in point—return only the Statistic Name along with a value.

If you create a *Delta* statistic called *HubMail.Transferred.Hourly*, for instance, it returns simply:

HubMail.Transferred.Hourly=[value]

Bear this in mind when creating Statistic Names.

The name of a Mail Outgoing Attachment Type should not *include the file extension, for instance, since PM adds it automatically when it returns data.*

The name of a Delta statistic, by contrast, should *include the reset interval, since PM doesn't add it when it returns data.*

Example 2: Generating data by type

We'll stick with the *Mail Outgoing Attachment Type* for now, and assume that, on a given server collecting this statistic (and for a given collection period), the following is true:

- Statistic Name: OutAttchExt
- Outgoing messages with attachments: 50
- Total volume of these messages: 10MB
- Attachment type: PDF only

PM would generate the following data:

OutAttchExt.pdf.Count=50
OutAttchExt.pdf.Size=10MB

You can *control the* extensions *for which this statistic type generates data—or opt to generate data for files of all types.*

You cannot*, however, selectively generate individual components of what a PM statistic returns (such as the* Count *or* Size *component, in the above example).*

For details on what is returned by all available PM statistic types, see *Data Returned by PM Statistic Types* starting on *page 459*.

5.1.4.0 Analyzer statistic definitions

Whereas PM statistics *generate* data, Analyzer statistic definitions enable you to *retrieve* data.

5.1.4.1 Essential differences between PM and Analyzer statistics

There are two fundamental differences between PM and Analyzer statistics:

- Data generation vs retrieval
  - PM statistics *generate* data.
  - Analyzer statistics *retrieve* data.
- Selective generation/retrieval
  - PM statistics generate *all* components of a given statistic type.
  - Analyzer statistic definitions can *selectively* retrieve data components.

Example 3: Retrieving selected data

Continuing with the previous example, let's assume you want to report on the *Count* of outgoing messages with attached PDF files, but not the *Size* (or volume in bytes) of those messages. What Analyzer statistic definition would accomplish that? Simple.

OutAttchExt.pdf.Count

(For information on how to create this definition, see *9.4.1.0 on page 299*.)

5.2.0.0 MANAGING DATA

5.2.1.0 Size of iwstats.nsf

Depending on the statistics you have enabled, and how replication is configured, your iwstats.nsf database—the repository for Performance Manager's statistics—can grow very large.

For example, in a 250-server environment with an iwstats.nsf of 3 GB, total disk usage would be (250 * 3 GB), or 750 GB, an unacceptable amount when taking into consideration the cost of managing disk arrays and the cost of backup—and then there's the cost in bandwidth during periodic replication of that database.

5.2.1.1 Limiting the growth of iwstats.nsf

We recommend that you configure your systems so that only one copy of iwstats.nsf contains data for all servers in your environment, namely the copy that resides on the Analyzer server.

In short, we suggest that you configure iwstats.nsf for one-way replication.

There are two ways to set up this one-way replication:

- Replication formulas
- ACL restrictions

We recommend the second method, since it's easier to implement and maintain.

5.2.1.2 Setting up the ACLs

The ACL entry for the Hub server itself (that is, where IntelliWatch Analyzer runs) should assign that server READER access, as in *Figure 5-1*, below.

FIGURE 5-1: ACL setting for the Analyzer Hub

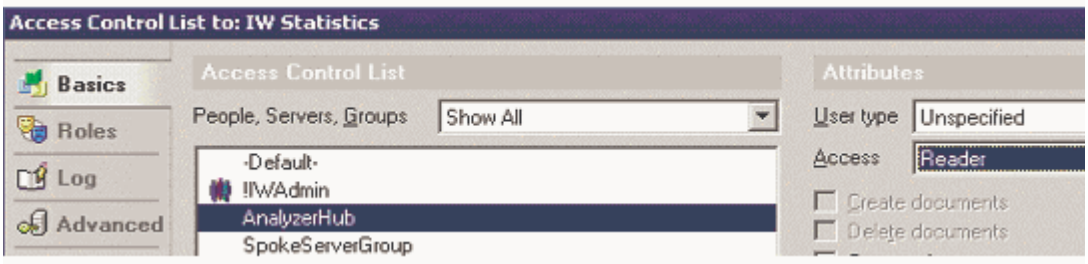

The ACL entry for the Spoke servers (where the data is being collected) should assign those servers DEPOSITOR access, as in *Figure 5-2*, below.

FIGURE 5-2: ACL setting for the Spoke Servers

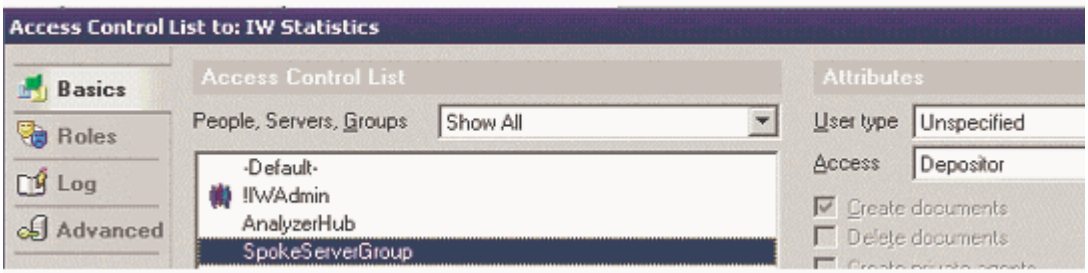

5.2.1.3 Picturing the data flow

The following flow chart illustrates the direction in which data is transferred.

FIGURE 5-3: Data flows from Spokes to Hub--but NOT in the other direction

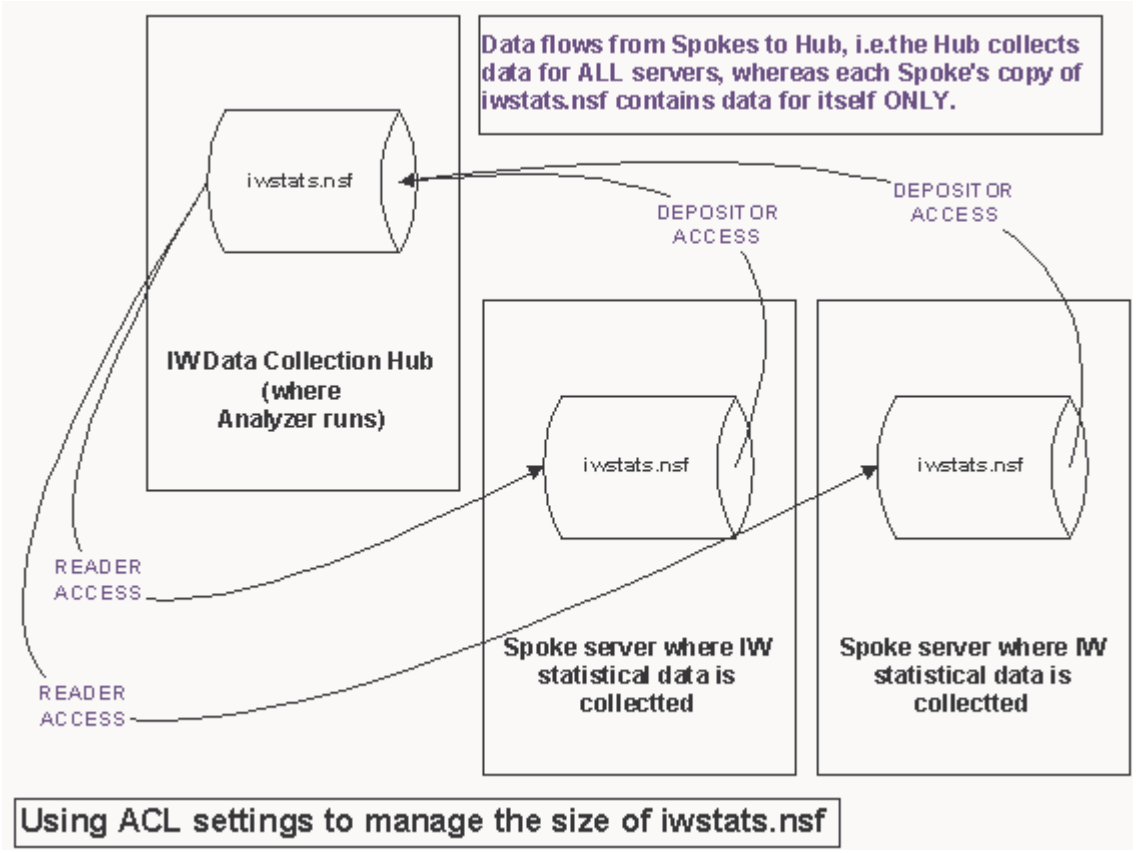

5.2.2.0 Long-term data management

Even using the above scheme to control the growth of iwstats.nsf on your servers, eventually you'll need to warehouse your IntelliWatch data. Use the database's Space Saver settings to accomplish this easily and quickly.

5.2.2.1 Space Saver settings

Open iwstats.nsf, then go to **File > Replication > Settings > Space Savers** to configure the Space Saver settings (for the dialog, see, below).

FIGURE 5-4: Space Saver Dialog

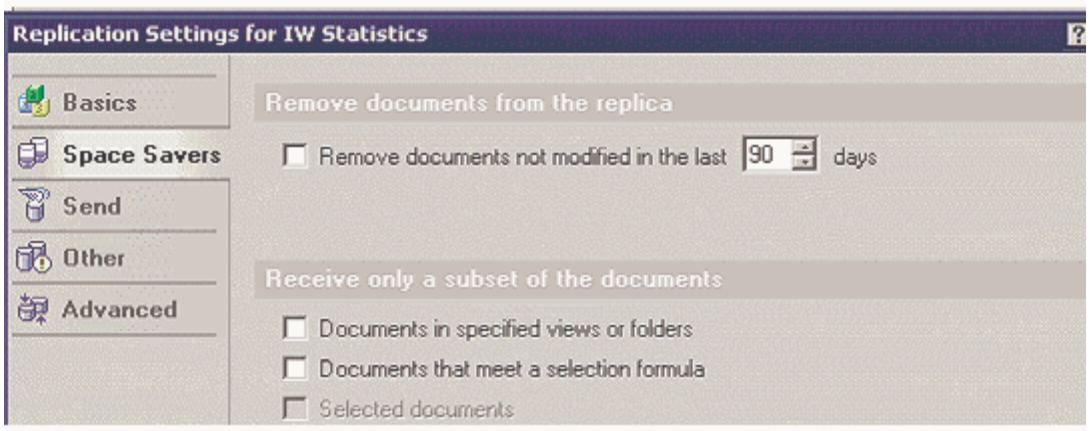

Three days should be adequate on the spokes; make the setting as long as needed on the Hub where Analyzer runs.

Note: Normally the Hub will hold 45 days of data, and will have a program in place to make a copy of the database every 30 days, so that individual months of data can be recovered as required.

5.3.0.0 ARCHITECTURE

FIGURE 5-5: PM's interaction with IntelliWatch Monitor and Analyzer

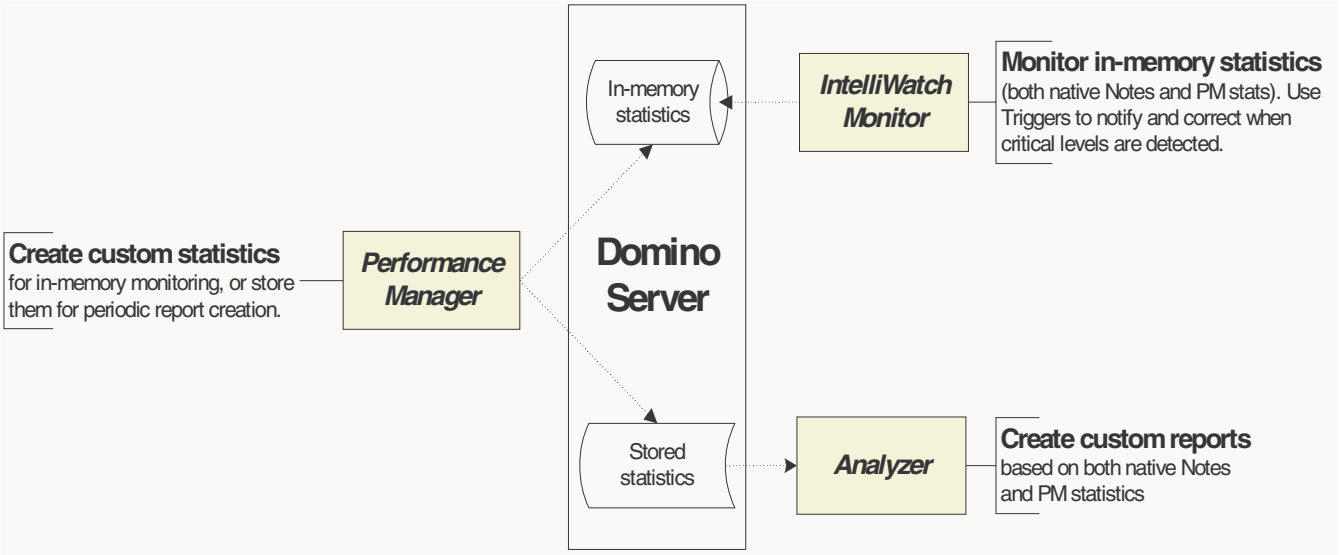

5.3.1.0 Interaction with Monitor

5.3.1.1 Triggers and in-memory stats

Triggers of the *Statistic* type monitor Notes and PM stats. (The default reset interval for these in-memory statistics is 60 minutes; the minimum allowed value is 5 minutes; the value you choose must be a multiple of 5.)

If the 'Enable default Triggers' option is selected during the Setup, one Statistics Trigger is enabled out-of-the box (in **iwserver.nsf**). This Trigger checks the value of the native Notes statistic **Server.Sessions.Dropped** (and initiates notification actions if the threshold is exceeded.)

In the same way, Statistic Triggers can monitor any IntelliWatch statistics you create, and take actions appropriate to the situation.

To make a PM statistic available for in-memory monitoring by Triggers, simply select the *Monitoring* option in the *Statistic Usage* section of the Monitoring Information tab, as illustrated in Figure 5-6, below.

FIGURE 5-6: Making statistics available to Monitor

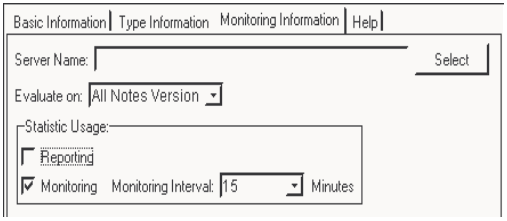

5.3.2.0 Interaction with Analyzer

5.3.2.1 Reporting on collected statistics

Out-of-the box, Pinnacle Performance Manager creates the Mail statistic **Delivery.Exceeds.OneHour**.

As soon as the Pinnacle Performance Manager add-in task is started on a Domino server, data for this statistic are collected and written to the IntelliWatch statistics repository, **iwstats.nsf**.

(By default, this statistic is also configured for in-memory monitoring by IntelliWatch Triggers. Both Reporting and Monitoring check boxes are selected, as can be seen in *Figure 5-7*.)

FIGURE 5-7: Making statistics available to IntelliWatch Analyzer

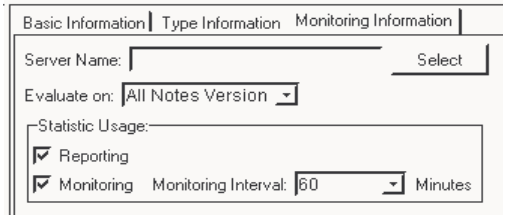

5.4.0.0 STATISTICS: CATEGORIES VS TYPES

5.4.1.0 Terminology

Before proceeding with statistic creation and configuration, you need to understand the terms *Category* and *Type* as they apply to IntelliWatch Statistics.

5.4.2.0 Category

Categories represent the folders displayed in the tree view in the left-hand pane of the PM user interface. They can help you organize statistics into meaningful groups.

5.4.2.1 How categories are created

When you create statistics, the name is entered in two parts:

- **Prefix (select or enter in combo box)**
- **Balance of name (enter in text box)**

The prefix becomes the category. If no such category exists, a folder by that name is created (and the statistic displayed under it); otherwise, the statistic is displayed under

the folder corresponding to the (pre-existing) category.

Example 4: Statistic categories

You want a separate category for mail statistics on your main Hub server, called **HubMailStats**.

When you create the first of these statistics, enter **HubMailStats** in the combo box. A folder by that name is created; the balance of the statistic name appears beneath it.

HubMailStats.Mail.Dead.Hourly is displayed as the folder **HubMailStats**, with **Mail.Dead.Hourly** listed beneath it (see *Figure 5-8*, below).

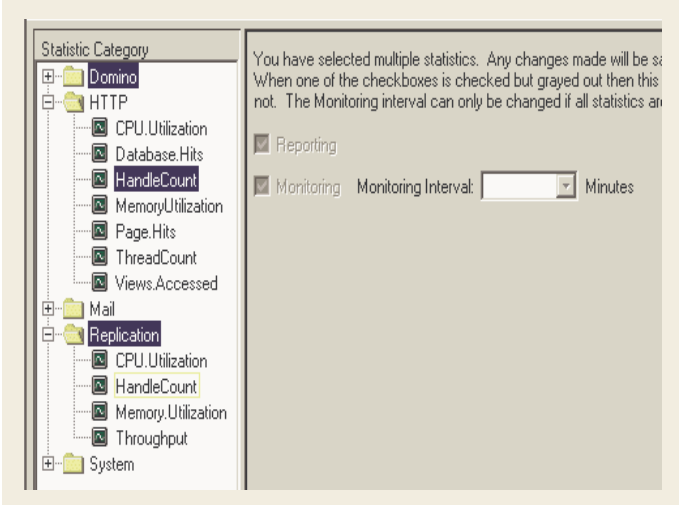

FIGURE 5-8: Category in tree view as folder

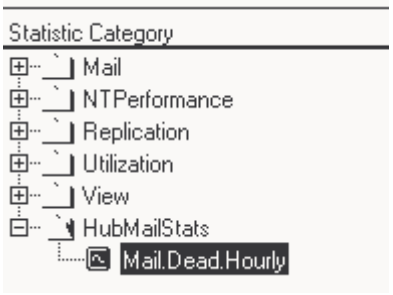

Additional statistics created in the category **HubMailStats** are added to the list beneath that folder.

Like IntelliWatch Monitor's Trigger categories, statistic categories are for your convenience only. They have no influence on functionality.

5.4.3.0 Type

Types represent the *nature of the counting being done*:

- Average
- Delta
- Difference
- File System Information
- Mail Incoming Delivery
- Mail Domain

- Mail Outgoing Attachment Types
- Mail Outgoing Server Attachment Percentage
- Mail Outgoing Server Volume
- Mail Size
- NT Performance Counter
- Replication
- Replication Delay
- Server Event Count
- Summation
- View Performance

The configuration options available for a given type are too complex to describe in brief. For details by type, see *page 218* to *page 249*, as well as *Data Returned by PM Statistic Types* starting on *page 459*.

5.5.0.0 CUSTOMIZING THE PM EDIT MENU

The drop-down Edit menu of the PM interface requires some explanation before moving on to the subject of statistic creation.

FIGURE 5-9: PM Edit menu

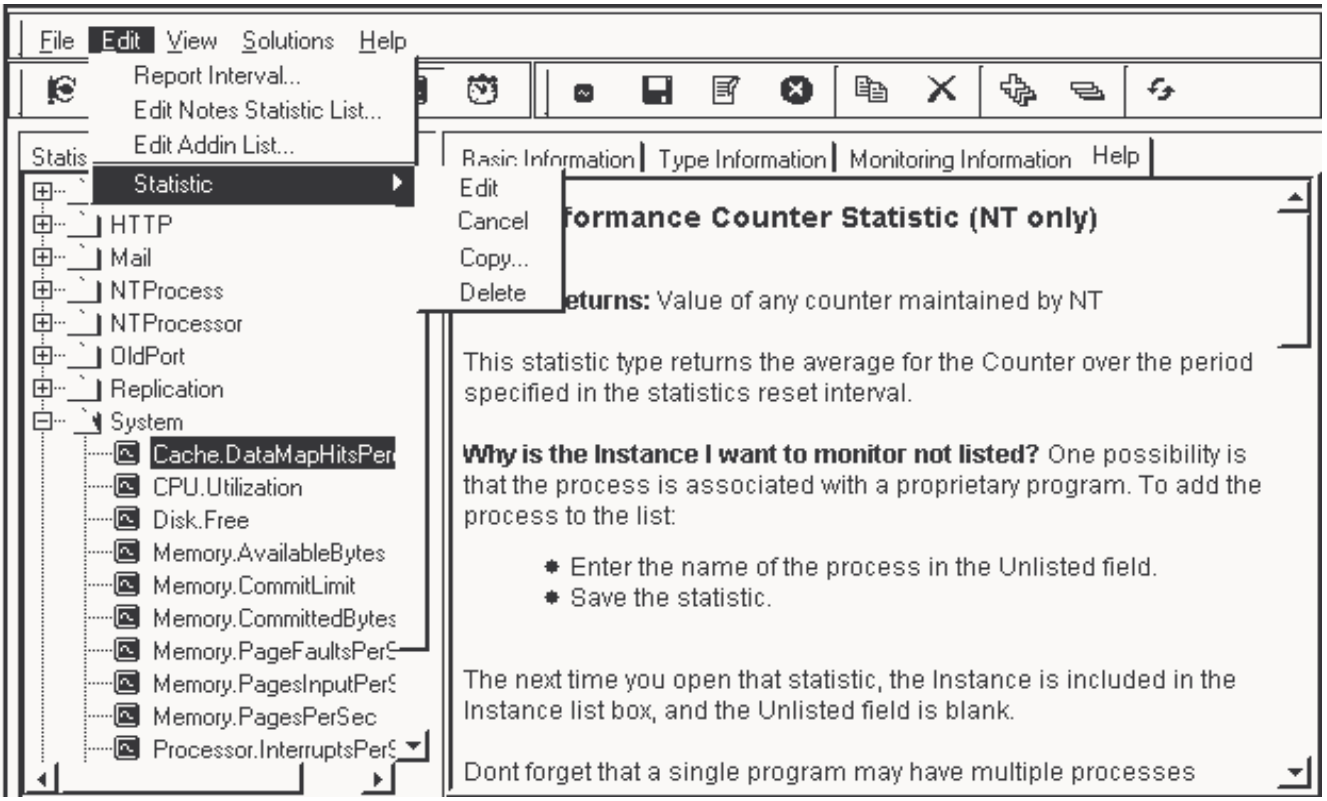

5.5.1.0 Report interval

This setting represents the frequency with which statistics are written to **iwstats.nsf**. The default value is 60 minutes.

Selecting this menu item brings up the dialog in *Figure 5-10*.

FIGURE 5-10: Configuring Report Interval

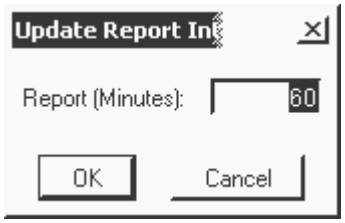

Setting the Report Interval to a value substantially below 60 minutes is not recommended, especially if a large number of statistics are being written to the database.

5.5.2.0 Edit Notes Statistic List

This customizable list of Notes native statistics is not intended to be exhaustive, but rather a quickly accessible list of those statistics that are particularly useful in creating PM statistics.

Selecting this menu item brings up the dialog in Figure 5-11. Use this dialog to customize the list of statistics to the requirements of your environment.

FIGURE 5-11: Edit Notes statistic list dialog

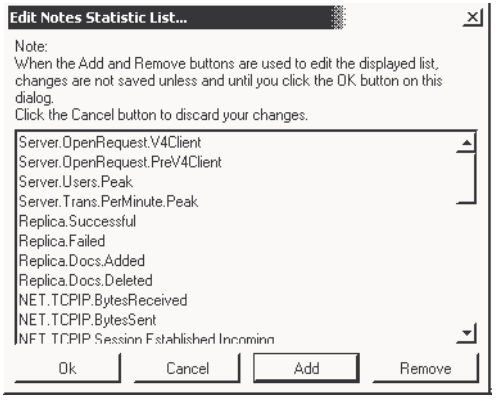

5.5.2.1 Adding statistics to the list

If you work frequently with statistics not displayed in the dialog, you may want to add them to the pop-up list.

TO ADD STATISTICS TO THE LIST:

1 Click the Add button. This brings up the dialog in Figure 5-12

FIGURE 5-12: Edit statistic list dialog.

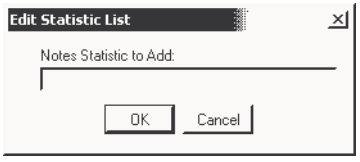

2 Type in the name of the Notes statistic you want to add.

Please type the statistic name accurately. No check is performed to confirm that what you typed in is actually a Notes native statistic. PM cannot use improper additions to the list.

3 Click OK on the dialog in Figure 5-12.

Though the displayed list now reflects the addition, if you do not click OK (in Step 4), no changes are made to the list on the server.

4 Click OK on the original dialog (see Figure 5-11) to confirm the addition.

5.5.2.2 Removing statistics from the list

The pop-up list may include statistics that are not relevant to your Domino environment. You may want to remove these statistics, to make the list more useful.

TO REMOVE STATISTICS FROM THE LIST:

1 Click on the statistic(s) you want to remove.

Shift and Ctrl keys function in the usual way (under Windows) to select multiple statistics.

2 Click the Remove button.

Though the displayed list now reflects the deletion, if you do not click OK (Step 3), no changes are made to the list on the server.

3 Click OK on the original dialog (see Figure 5-11) to confirm the deletion.

5.5.3.0 Edit Add-in List

This customizable list of Notes add-in tasks is not intended to be exhaustive, but rather a quickly accessible list of those add-in tasks that are particularly useful in creating PM statistics.

Selecting this menu item brings up the dialog box (in Figure 5-13), allowing you to add server tasks to the list, or remove them.

FIGURE 5-13: Edit add-in list dialog

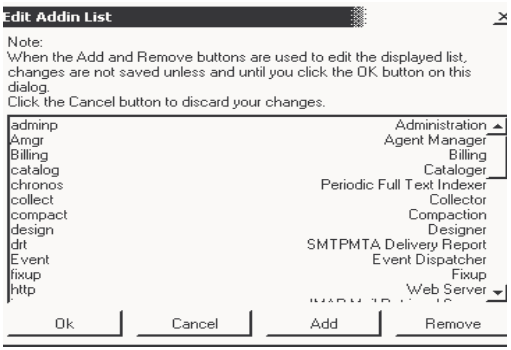

5.5.3.1 Adding new tasks to the list

If you work frequently with add-in tasks not displayed in the dialog, you may want to add them to the pop-up list.

TO ADD ADD-IN TASKS TO THE LIST:

1 Click the Add button.

This brings up the dialog in Figure 5-14

FIGURE 5-14: Edit add-in list dialog

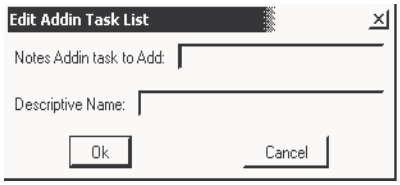

2 Type in the name of the Notes add-in task you want to add to the list.

After you complete Step 4, below, the task is displayed on the left-hand side of the list box in Figure 5-13.

Please type the add-in name accurately. No check is performed to confirm that new entries are actually Notes add-in tasks. PM cannot use improper additions to the list.

3 Type in a meaningful Descriptive name.

This name is for your convenience only, and has no effect on functionality.

4 Click OK on the dialog in Figure 5-14.

Though the displayed list now reflects the additions, if you do not click OK (in Step 5), no changes are made to the list on the server.

5 Click OK on the original dialog (Figure 5-13) to confirm the addition.

5.5.3.2 Removing add-ins from the list

The pop-up list may include add-ins that are not relevant to your Domino environment. You may want to remove these statistics, to make the list more useful.

TO REMOVE ADD-INS FROM THE LIST:

1 Click on the add-in(s) you want to remove.

Shift and Ctrl keys function in the usual way (under Windows) to select multiple items.

2 Click the Remove button.

Though the displayed list now reflects the deletion, if you do not click OK (in Step 3), no changes are made to the list on the server.

3 Click OK on the original dialog in (see Figure 5-13) to confirm the deletion.

5.5.4.0 Edit > Statistic...

Use this menu item to access the standard Edit, Cancel, Copy, and Delete options, which apply to the currently selected statistic.

Once a statistic is selected—*but before you put it into Edit mode*—Edit, Copy and Delete are enabled, but not Cancel. Once in Edit mode, only Cancel is enabled.

5.5.5.0 Reporting vs Monitoring

As illustrated in *Figure 5-5 on page 196*, PM statistics can be configured both for use in report creation and for in-memory monitoring.

These two options are configured on the Monitoring Information tab of the PM statistic template (see *"Monitoring Information" on page 215*).

5.5.5.1 Reporting

- If selected, a statistic is written to **iwstats.nsf**, for later use by Analyzer.

To set the interval at which the statistic is written to the database, go to **Edit > Report Interval** via the drop-down menus (see *Figure 5-9 on page 200*).

The Report Interval setting is global, that is, it cannot be configured for individual statistics.

This setting has a lower limit of 15 minutes, and must be a multiple of 5.

5.5.5.2 Monitoring

- If selected, the statistic is written to memory, for use by Monitor Triggers.

Unlike the Report Interval, the Monitoring Interval (at which in-memory statistics are recalculated and updated) can be configured for individual statistics. This is done on the

The Monitoring Interval has a lower limit of 5 minutes, and must be a multiple of 5.

This interval can be changed only if the Monitoring checkbox is selected.

5.5.6.0 Configuring multiple statistics

Several statistic features can be enabled/configured at the Pinnacle Console for multiple statistics.

- Enable Reporting
- Enable Monitoring
- Set Monitoring Interval for multiple statistics

5.5.6.1 Multiple selection procedure

To selectively configure PM statistics, simply Ctrl+click on the desired individual statistics and/or folders. (This causes the dialog in *Figure 5-15* to be displayed.)

FIGURE 5-15: Multiple-configuration dialog

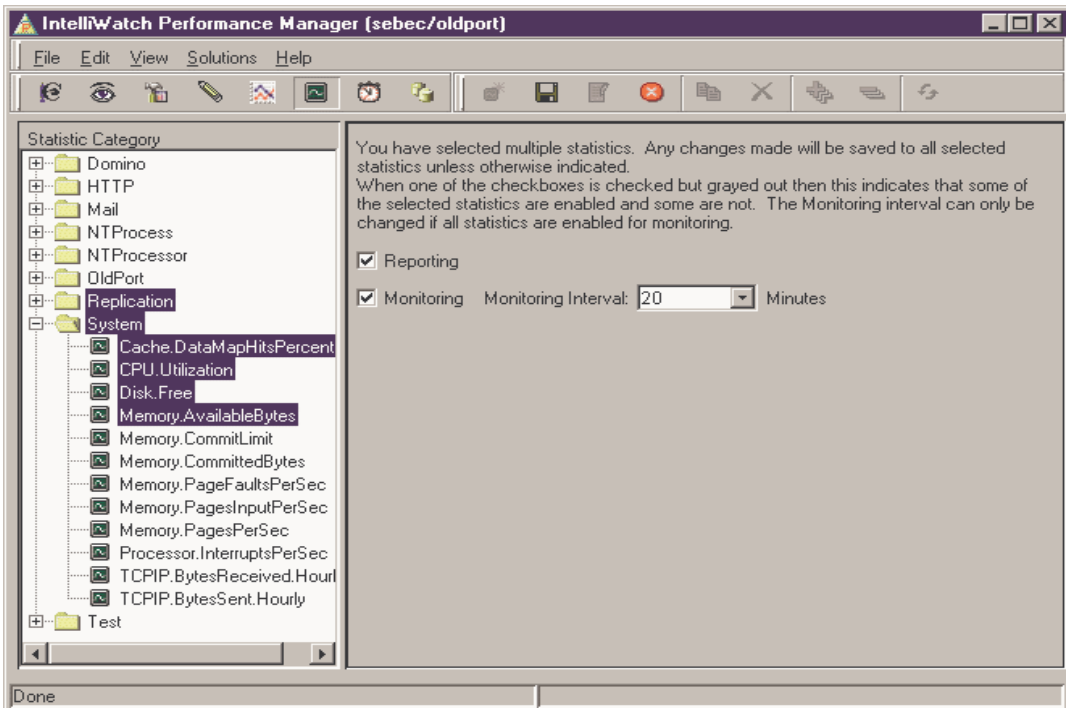

Checkboxes are triple-state: a gray check on a darker-gray background indicates only some of the selected statistics are configured for Reporting/Monitoring (see Figure 5-15.)

To edit these settings, click the  icon, or go to **Edit > Statistic > Edit** via the drop-down menus. Please see the following list of possible actions, along with their consequences.

- Click Reporting
  - Enables ALL selected statistics for Reporting (at the currently configured global Reporting Interval).

- Click Monitoring
  - Enables ALL selected statistics for Monitoring.
  - Changes Monitoring Interval to 15 (minutes)
- Change Monitoring Interval
  - New value will be applied to all selected statistics—but NOT to any other statistics currently enabled for Monitoring.

In the case of the selections displayed in *Figure 5-15*, the following statistics will be enabled for both Reporting and Monitoring:

- all statistics in the Replication folder
- selected statistics in the System folder

No other statistics will be affected. If they are already enabled, they will remain so; their Monitoring Interval will not change.

5.6.0.0 TYPE VS MONITORING INFORMATION

5.6.1.0 Mail Statistics

Certain PM Mail statistics have server lists on both the Type Information and the Monitoring Information tabs. Why is this?

5.6.1.1 Mail Statistic Server Lists

To illustrate the function of the two server lists, let's use the following example.

Example 5: Mail Incoming Delivery

- PM Type: Mail Incoming Delivery
- Statistic Name: Hub.InMail
- Reporting: enabled
- Monitoring: disabled
- Type Information server list:
  - ABC/NE/MailSvr1
  - DEF/NE/MailSvr2
  - GHI/West/MailSvr1
- Monitoring Information server list:
  - CentralOffice/MailHub1
- Options: CountOver ONLY

For incoming mail, this statistic type returns:

[Statistic Name].[server name].CountOver

and/or

[Statistic Name].[server name].Max
[Statistic Name].[server name].Avg
[Statistic Name].[server name].Count

The [server name] element refers to the originating server of qualifying mail traffic.

Assuming there is mail traffic from all three servers (on the Type Information tab) that meets the CountOver condition (see *page 229* for details), the following statistical data will be written to **iwstats.nsf** *on the CentralOffice/MailHub1 ONLY* (the only server listed on the Monitoring Information tab)*:*

Hub.InMail.ABC/NE/MailSvr1.CountOver=[value]
Hub.InMail.DEF/NE/MailSvr2.CountOver=[value]
Hub.InMail.GHI/West/MailSvr1.CountOver=[value]

(See *Figure 5-16*, below.)

5.6.1.2 Monitoring Server

Put simply, the Monitoring Information tab governs the server(s) to whose local copy of

iwstats.nsf data is written (assuming the statistic is enabled for Reporting), or ON which the statistic will be created in memory (assuming the statistic is enabled for Monitoring).

5.6.1.3 Type Information Server List

Put simply, the Type Information tab governs the server(s) whose mail traffic causes the Monitoring Server to generate data, either by writing it to the IntelliWatch statistics database, or by creating it in memory.

FIGURE 5-16: Server Lists for certain PM Mail statistics

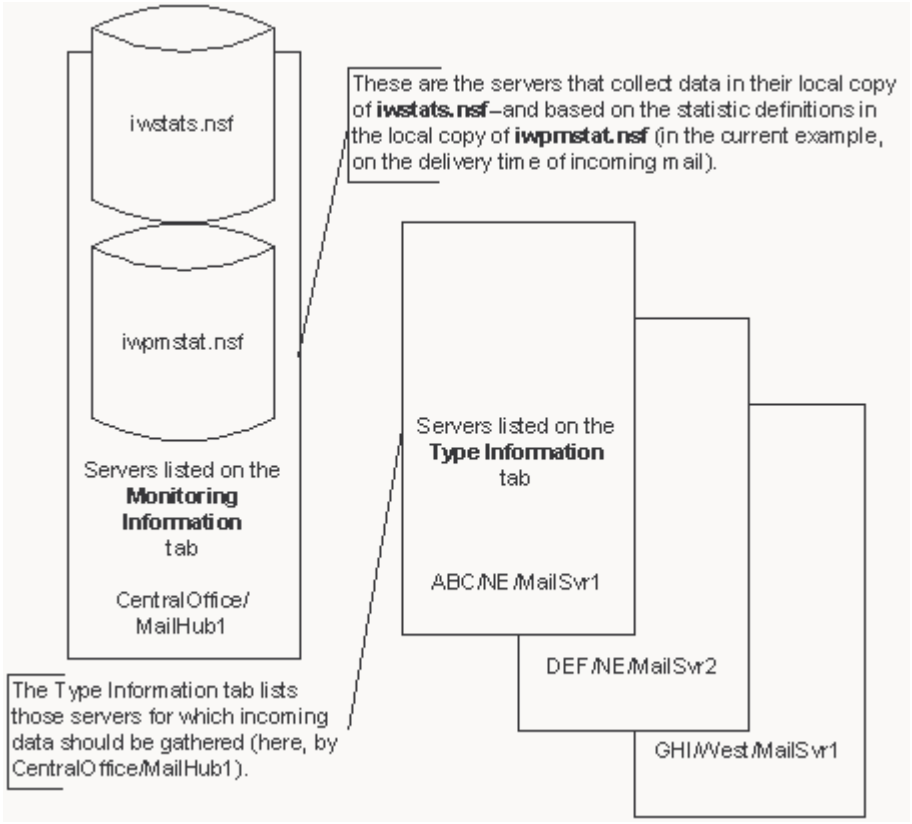

5.7.0.0 WORKING WITH STATISTICS

5.7.1.0 PM user interface

The PM user interface is comprised of three tabs:

- Basic Information
- Type Information
- Monitoring Information.

While it's possible to set out an unvarying set of steps for the first and last of those tabs, Type Information differs too widely from one PM statistic type to another to permit of one set of instructions. Type Information tabs will be discussed individually, starting on *page 219*.

The following example is provided to help you familiarize yourself with the user interface.

Example 6: Server Event Count

Table 5-1. Statistic requirements

Parameter	Description
Name prefix (category)	SvrEvtCts
Name suffix	iwaswDBAccess.ASWHubGroup
Type	Server Event Count
Servers	ASW_HubGroup
Description (purpose)	Data collection for report generation.
Notes versions	all versions
Database	iwasw.nsf
Reporting/ Monitoring	Report only
Server Event	DBAccess Open
Add-in task	IWASW

TO CREATE THE STATISTIC SVREVTCTS.IWASWDBACCESS.ASWHUBGROUP:

5.7.1.1 Basic Information

1 Select **File > New** from the menu bar; alternatively, click on the  toolbar icon.

2 Enter SvrEvtCts in the combo-box section of the Name field.

3 Enter iwaswDBAccess.ASWHubGroup in the text box of the Name field.

4 Select Server Event Count from the Statistic Type combo box.

5 Enter the Description from the table *"Chapter 5, Statistic requirements"*, above.

6 Switch to the Type Information tab.

TYPE INFORMATION

1 In the Internal Server Events section of the tab, select DBAccess.
2 Click the Lookup button to the right of the Add-in task text field.
3 Select iwasw in the left-hand list box, then move it to the Selected list using the Add button.
4 Click OK to return to the Type Information tab.

No User Name needs be specified in this instance.

5 Click the Lookup button to the right of the Database text field.
6 Select **iwasw.nsf** in the left-hand list box, then move it to the Selected list using the Add button.
7 Click OK to return to the Type Information tab.
8 Switch to the Monitoring Information tab.

5.7.1.2 Monitoring Information

1 Click the Lookup button to the right of the Server Name text field.
2 Select ASW_HubGroup in the left-hand list box, then move it to the Selected list using the Add button.
3 Click OK to return to the Monitoring Information tab.
4 Using the combo box, make sure All Notes Versions is the selected option.
5 Select the checkbox that activates Reporting. (If Monitoring is selected, deselect it.)
6 Save the statistic definition by going to **File > Save** via the drop-down menus; alternatively, click on the toolbar icon.

5.7.1.3 Help

The Help tab provides type-specific information. Displayed information cannot be edited.

5.7.2.0 Editing vs Copying statistics

Before you edit a statistic, consider if you'll want to resort to it later in its current form.

If the answer is 'yes', make a copy of the statistic, *and edit the copy*.

If, on the other hand, you are changing configuration details because they no longer apply, editing is probably the best course of action. In that case, follow these steps:

A statistic's Name field cannot be edited.

If all you want to do is change the name of a statistic, make a copy of it, then change the name of the copy. To avoid confusion, you may want to delete the original statistic.

1 Open the appropriate Statistic Category folder and select the statistic you want to edit.
2 Go to **Edit > Statistic > Edit** via the drop-down menus; alternatively, click on the toolbar icon.

The information in the right-hand pane now becomes active.

3 Make the required changes.
4 Save your changes by going to **File > Save** via the drop-down menus; alternatively, click on the toolbar icon.
5 Click OK to confirm the changes.

5.7.2.1 Copying statistics

Copying statistics is useful when you need to create several statistics that only vary in one or two parameters. You need fill in all

the required template fields just once, making the necessary adjustments to perhaps only one or two fields of the new copy.

Give some thought to the Name of the new copy. Make it reflect as closely as possible 1) what's being collected, and 2) from where.

Remember that two statistics cannot have the same Name.

TO COPY A STATISTIC:

1 Select the statistic you want to copy.

If the relevant folder is not expanded, click on the plus sign (+) to the left of the folder.

2 Go to **Edit > Statistic> Copy** via the drop-down menus; alternatively, click on the  toolbar icon.

3 Alter the Name in a way that reflects the role of the new statistic.

4 To save your changes, go to **File > Save** via the drop-down menus; alternatively, click on the  toolbar icon.

5 Click OK to confirm the action.

5.7.3.0 Deleting vs Deactivating statistics

You can stop the collection of a PM Statistic on a server in two ways:

- delete the statistic from **iwpmstat.nsf**
  - Use this method *only* if you are *sure* you no longer need to collect/monitor this statistic *on any of your servers*.

Since this database normally replicates to all servers running PM, deletions are propagated to all copies of the database.

- deselect the server from the list on which the statistic is to be created
  - Preferred method of stopping collection/monitoring of this statistic.

TO DELETE STATISTICS ON ALL SERVERS:

1 Select the statistic you want to delete, then go to **Edit > Statistic > Delete** via the drop-down menus; alternatively, click on the  toolbar icon.

2 Click OK to confirm the deletion.

3 Replicate the edited copy of **iwpmstat.nsf** to all servers running IWSTATG.

5.7.3.1 Deactivating statistics

Deactivation is preferable to deletion. Statistics can quickly be made available again, without your having to recreate them.

TO DEACTIVATE A STATISTIC ON SELECTED SERVERS:

1 Select the statistic you want to deactivate, then go to **Edit > Statistic > Edit** via the drop-down menus; alternatively, click on the  toolbar icon.

You are now in edit mode.

2 Go to the Monitoring Information tab, and deselect the servers/groups you no longer want to collect/monitor the statistic.

3 To save your changes, go to **File > Save** via the drop-down menus; alternatively, click on the  toolbar icon.

4 Replicate the edited copy of **iwpmstat.nsf** to all servers running IWSTATG.

TO REACTIVATE THE STATISTIC ON THE AFFECTED SERVERS:

1 Select the statistic you want to reactivate, then go to **Edit > Statistic > Edit** via the drop-down menus; alternatively, click on the  toolbar icon.

You are now in edit mode.

2 Go to the Monitoring Information tab, and select all servers/groups on which the statistic is to be collected/monitored.

3 To save your changes, go to **File > Save** via the drop-down menus; alternatively, click on the  toolbar icon.

4 Replicate the edited copy of **iwpmstat.nsf** to all servers running IWSTATG.

5.8.0.0 COMMON STATISTIC FIELDS

The Basic Information and Monitoring Information tabs are the same for all statistic types. For field definitions and usage suggestions, see *page 212* through *page 215*.

5.9.0.0 STATISTIC FIELDS BY TYPE

The fields of the Type Information tab vary by statistic type. For field definitions and usage information for a specific type, see *page 218* through *page 249* (arranged alphabetically).

PM: Basic Information

Basics ...

Use this dialog to:

- select a statistic name
  - two-part names allow flexibility in categorizing

    The *Prefix* (in combo box) becomes a folder for categorizing stats.

    The *Suffix* (in text box) is listed under the folder represented by the Name prefix
- select a statistic type
  - represents type of data collection/manipulation
- enter an optional description
  - information about statistic's purpose in your environment

Putting it into practice ...

Naming your statistics:

Name statistics in ways meaningful to you, and representative of what data are being collected (and from where). Just observe the following restrictions:

- only valid characters may be used
  - letters, numbers, underscore, dollar sign
  - spaces are *not* allowed
- names must begin with a letter
- names may not exceed 255 characters (limitation imposed by Notes)

Making the best use of the Basic Information tab:

- Create Name prefixes (*Categories)* that allow you to see *at a glance* the function of a group of statistics.
  - If you have several statistics relative to Mail.Dead, for instance, create a separate *DeadMail* category, instead of putting these statistics in the *Mail* folder.
- Use the *Description field* to record information not readily apparent. Someone else may later have to work with this statistic.
  - Avoid stating the obvious: it obscures the salient points of your description.

PM Common Dialog: Basic Information

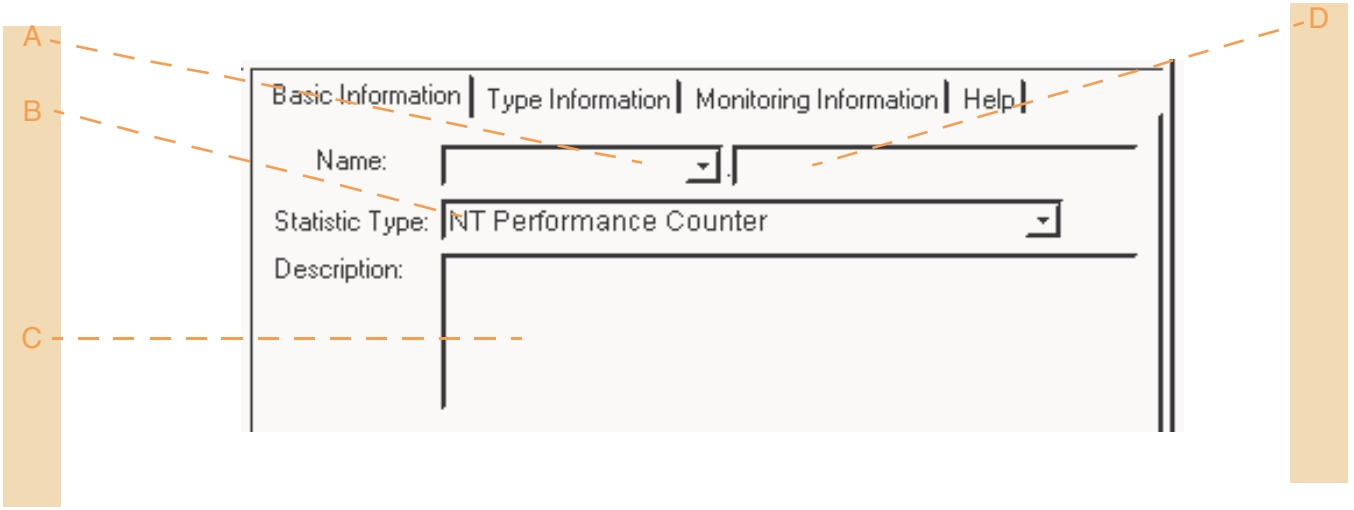

A: statistic prefix (category). Balance of name at E.
B: statistic type
C: statistic's purpose in your environment
D: balance of statistic name

PM: Monitoring Information Tab

Basics ...

Use this dialog to:

- configure on which servers to create a statistic
  - click the Select button to bring up a list box containing servers and server groups (For details on this server list, see *"Monitoring Server" on page 206*.)
- select a Notes version (or configure statistic as version-independent)
  - options are: Notes 4 only, Notes 5 only, All Notes Versions
- check Reporting to write statistic to iwstats.nsf
  - makes data available for reporting using IntelliWatch Analyzer
- check Monitoring to create in-memory statistic
  - makes statistic available for real-time monitoring using IntelliWatch Triggers
- enter interval at which to refresh in-memory stats

Putting it into practice ...

Making the best use of the Monitoring Information tab:

- conserve system resources by not collecting data you won't use
  - When selecting the servers/server groups on which to collect a statistic, take the extra few minutes required to narrow down the list to only those servers/groups for which the data are truly relevant.
- use the Notes version option to your advantage
  - Consider whether the statistic you are collecting/monitoring exists on both Notes 4 and R5 (or, if it exists, if it's significant on all Notes versions in your environment).
- for in-memory statistics, tailor the monitoring interval to the seriousness of exceeded thresholds

PM Common Dialog: Monitoring Information

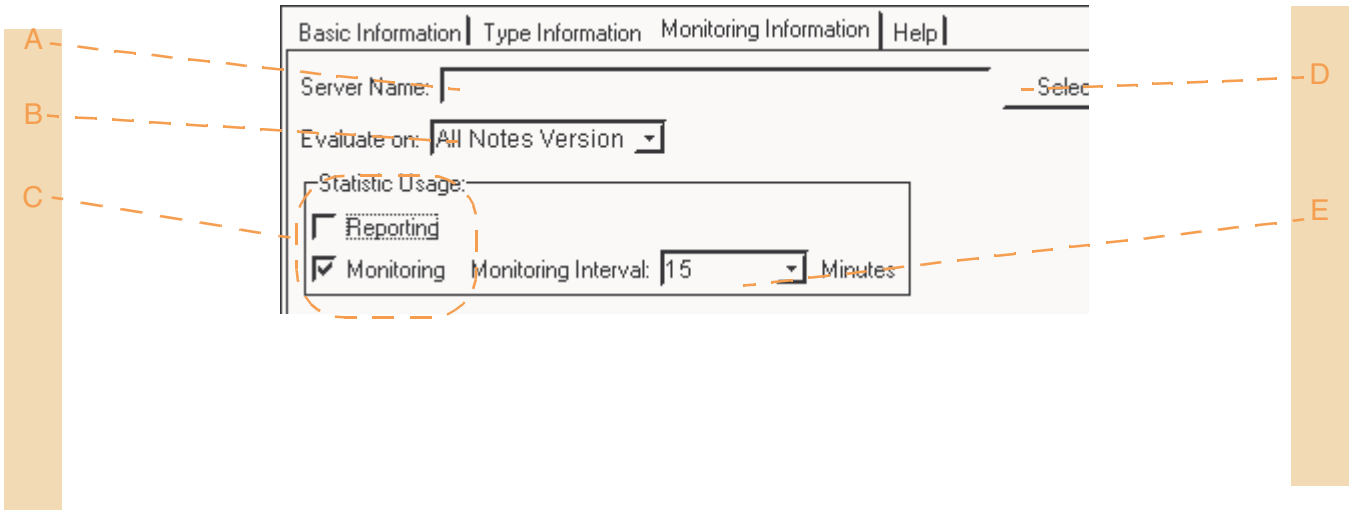

A: servers on which to evaluate statistic
B: Notes versions on which to evaluate statistic (R4 only, R5 only, or All)
C: select Reporting (for use with Analyzer), Monitoring (as in-memory statistic), or both
D: button to launch Select Server dialog
E: time interval for monitoring

PM: Help Tab

Basics ...

Use this dialog to:

- obtain type-specific usage information

Putting it into practice ...

Please note:

The Help text provided is not editable.

PM Common Dialog: Help

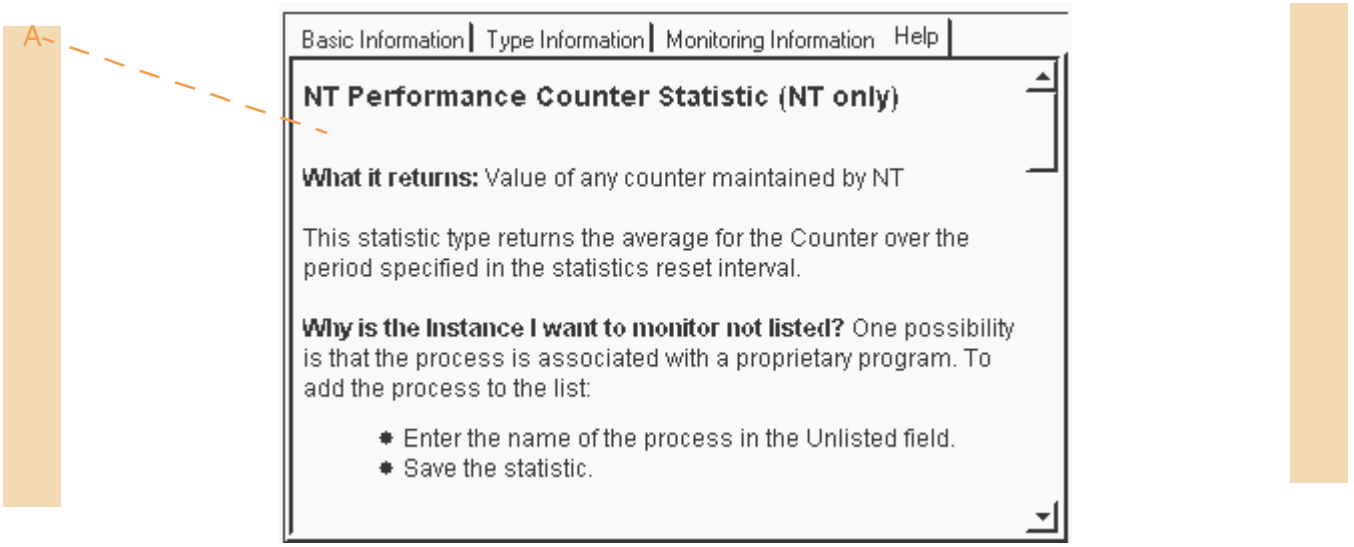

A: non-editable, type-specific usage information

PM Statistic Types: Average

Basics ...

- reports Average of two or more statistics

A list of commonly averaged statistics is created by the Setup, and displayed on the Type Information tab.

As the word Average implies, *at least two* statistics must be selected at the Type Information tab. If only one statistic is selected, the save operation fails.

Adding unlisted statistics:

If a statistic you want to average does not appear in the drop-down list, you have two options:

- type it in, then save the statistic

  All other required fields must be filled in first, or the save operation fails.

- go to Edit > Edit Notes Statistic List via the drop-down menus to add statistics

  For details on this procedure, see *"Edit Notes Statistic List" on page 201*.

Using the first method makes the added statistic available for the current statistic only.

Using the latter method makes the addition available to all statistics—even those already in existence.

Putting it into practice ...

Example:

Geological Services Company wanted to establish a baseline average for CPU usage by Domino server tasks.

Firstly, they created NT Performance Counter statistics for each of the tasks running on their system, from Router to HTTP. (These statistics were configured to check the *% Processor Time* on the relevant Process object, and were named succinctly as UsageCPU.[TaskName].)

Secondly, they created a statistic of the Average type to process the values of all statistics created in the previous step.

To compare the average with the CPU usage of individual server tasks, Difference statistics were used.

PPM: Average

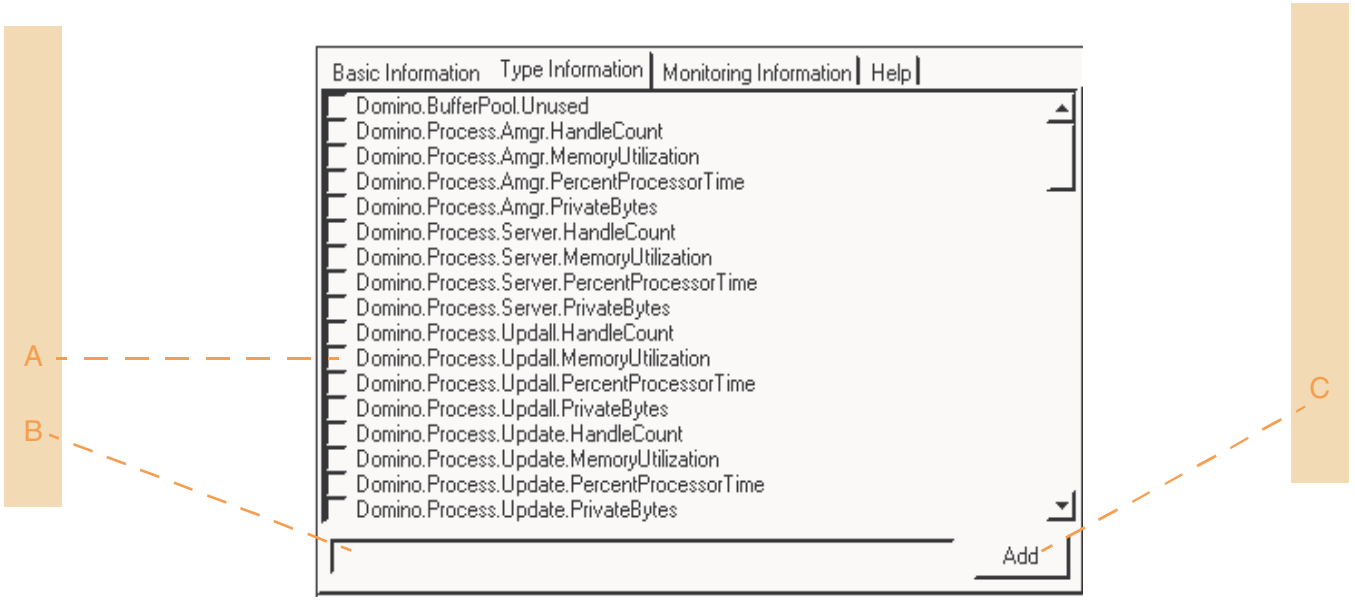

A: select statistics to be averaged
B: use to enter statistics not in the list (for usage, see previous page)
C: clicking the Add button both *adds* the manually entered statistic to list, and *selects* it

PM Statistic Types: Delta

Basics ...

- change in the value of another statistic over a user-specified time interval (two such statistics are discussed in the example, below)
  - Delta statistics have only one unique parameter, the name of the statistic to be reset.
  - The Reset interval is configured on the Monitoring Information tab.

The recommended reset interval is the same as the monitoring frequency of any IntelliWatch Triggers that are checking the value of this statistic.

A list of statistics commonly used to build Delta statistics is created by the Setup. This list is displayed on the statistic's Type Information tab.

If a statistic you want to use is not listed:

For information, see *"Adding unlisted statistics:" on page 218*, and *"Edit Notes Statistic List" on page 201*.

Putting it into practice ...

Example:

Evans Electrical Supply, Inc. needed to establish baseline values for NET.TCPIP.BytesReceived and NET.TCPIP.BytesSent, based on the time of day.

Delta statistics were created to track these two values on an hourly basis (monitoring internal=60 minutes).

Both the Monitoring and the Reporting options were selected, which made the data available for both monitoring by Triggers and report generation by IntelliWatch Analyzer.

PPM: Delta

A

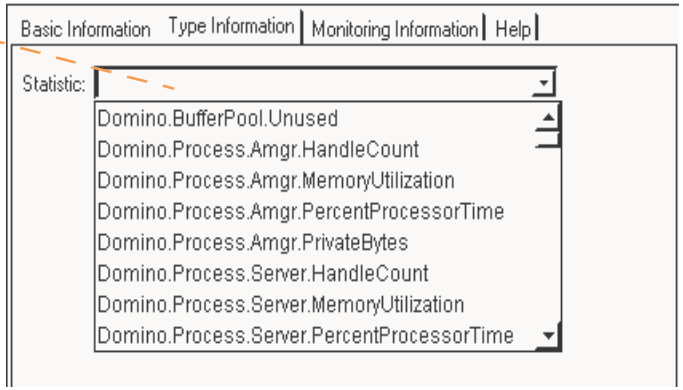

A: select statistic whose Delta value is to be computed

PM Statistic Types: Difference

Basics ...

Reports:

- difference between two statistics

A list of statistics is created by the Setup for which Difference is particularly useful. This list is displayed on the statistic's Type Information tab.

As the with the Average type, if only one statistic is selected, the save operation fails.

If a statistic you want to use is not listed:

For information, see *"Adding unlisted statistics:" on page 218*, and *"Edit Notes Statistic List" on page 201*.

Only useful if comparing apples and apples?

While Difference statistics are easier to interpret when both statistics are in the same units of measurement (see Example 1, below), they can also be useful when the statistics being compared are measuring 'apples and oranges' (see Example 2, below).

Putting it into practice ...

Example: 1:

A large difference between a mail server's *BytesReceived* and *BytesSent* statistic values can indicate that a significant amount of mail is being resent. This situation can adversely impact server performance.

Using the PM Difference statistic type, compare these two Notes statistics, and take appropriate notification/other actions, in the event the threshold value is exceeded.

Example: 2:

Answering the question 'Is *Server.Users.Peak* too high on Server XYZ?' can be easier if you know the value of *Server.Sessions.Dropped*.

A useful diagnostic ploy which adds meaning to both statistics involves creating a PM Difference statistic that compares *Server.Users.Peak* to *Server.Sessions.Dropped*.

Especially over time, this kind of comparison can assist Admins with load-balancing.

A system where a substantial increase in *Server.Users.Peak* is accompanied by a concomitant rise in *Server.Sessions.Dropped* is in greater need of additional resources than a system where substantially increased user activity is not accompanied by a similarly large jump in *Server.Users.Peak*.

PPM: Difference

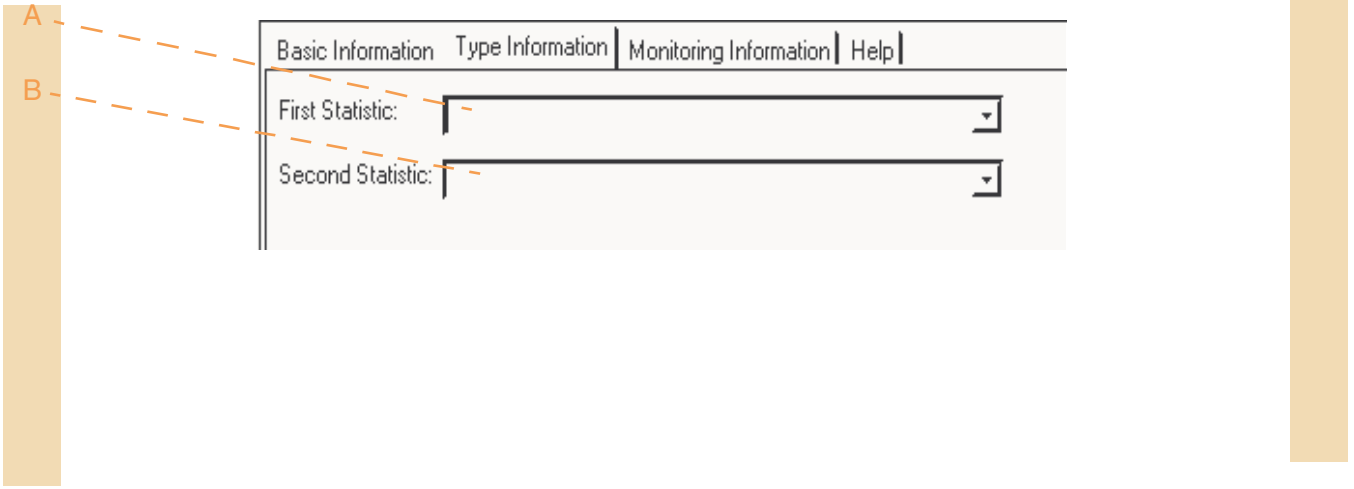

A: statistic to be subtracted from
B: statistic whose value is subtracted from the First Statistic

PM Statistic Types: File System Information

Basics ...

Reports (UNIX platforms only):

For each mount location:

- <statname>.<mount point>.Free
- <statname>.<mount point>.Size

This statistic type was designed as a UNIX counterpart to the Disk.[N].Free Notes statistic on NT, and has no relevance on the latter platform.

Putting it into practice ...

IncludedPartitionTypes:

If you encounter an error such as

"Warning: File system "/etc/mnttab" has an unknown type "mntfs" and will not be reported on",

edit the *IncludedPartitionTypes* setting in **iwperfmn.ini** to include the requisite types.

Typical contents of iwperfmn.ini:

[Stat Generator]

Database=iwstats.nsf

LogLevel=1

LogLevelConsole=0

TransactionBorderThreshold=40000

IncludedPartitionTypes=nfs,jfs,nfs3

ExcludedPartitionTypes=cdrfs

The File System Information statistic type requires no type-specific user input.

The statistic returns the same information (see previous page) for each non-NT server listed on the Monitoring Information tab (or for all known servers, if the field is left blank).

To configure this statistic type for Monitoring and/or Reporting, follow the same procedures as for all other IntelliWatch statistic types.

PM Statistic Types: Mail.Domain

Basics ...

Reports:

- Count (cumulative)

<Statname>.incoming.<domainname>.[count]

<Statname>.outgoing.<domainname>.[count]

- Total size in bytes

<Statname>.incoming.<domainname>.[size]

<Statname>.outgoing.<domainname>.[size]

A separate statistic is reported for each address in each domain fitting a 'pattern', both incoming and outgoing.

Patterns:

Patterns allow use of three wildcards: # (for any numeric character), ? (for any *single* character) and * (for zero or more characters).

Examples: *.com, *.org, t??m.org (which would collect statistics for both *Team.org* and *Term.org*).

Thresholds:

The Size and Count thresholds apply to in-memory statistics only, and have no effect on what is reported to the IntelliWatch statistics repository **iwstats.nsf**.

Putting it into practice ...

Count/Size thresholds in practice:

Count.threshold=200; Size.threshold=1MB

If 199 messages totaling 999KB are received from the domains listed, neither Count nor Size are available for in-memory monitoring. If reporting is enabled, these values *are* written to **iwstats.nsf**, regardless of their size.

If 201 messages totaling 500KB were received from the monitored domains(s), the Count value *would* be available for monitoring, although the Size value would not.

Please bear in mind, however, that a Statistic Trigger checking for a value in bytes (KBs, MBs, and so on), won't fire based on this condition, since the monitored parameter is Size, not Count, whereas the Count value is the only one available in memory.

Example:

Smith Enterprises regularly needed to assess the amount of mail traffic being generated by/received from several different Internet domains.

Using wildcards creatively (see under Patterns, above) they made statistic generation as broad as was required.

*When wildcards are used to generate data—especially true with PM's Mail types—such a large number of unique statistics can be generated that performance (and even functionality) are affected. (Just such a pattern is *.com.)*

PPM: Mail.Domain

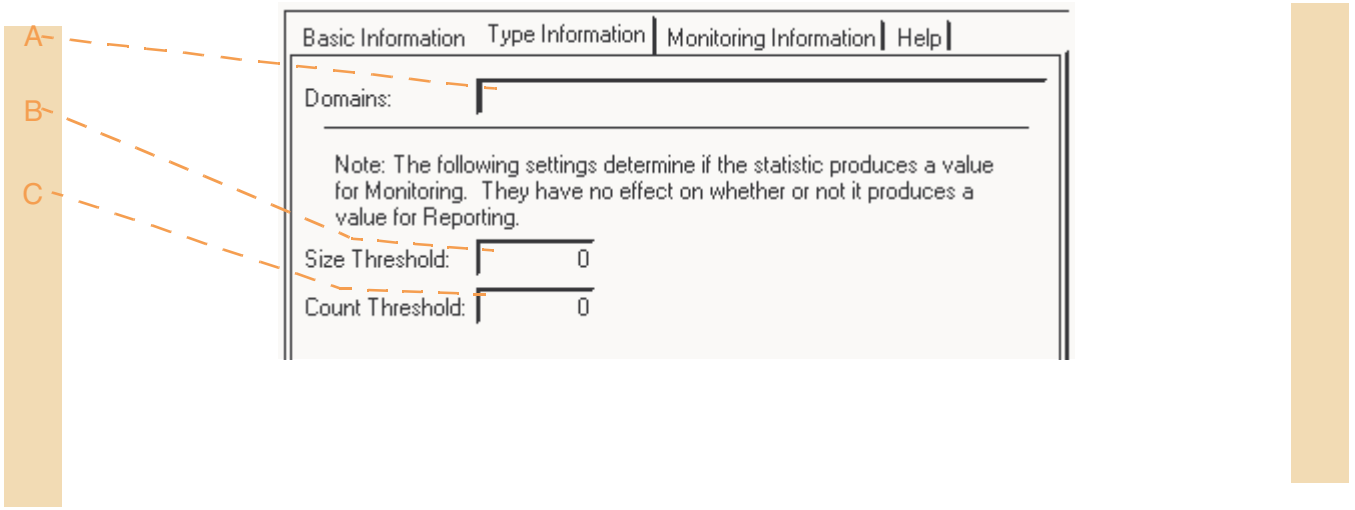

A: comma-delimited list of Domain(s) for which incoming mail to be checked
B: size threshold of incoming mail (applies to Triggers *only*)
C: count threshold of incoming mail (applies to Triggers *only*)

PM Statistic Types: Mail.Incoming.Delivery

Basics ...

Reports:

- number of times mail delivery time exceeded threshold, and/or,
- Min[imum], Max[imum] and Average delivery times
- broken down by:
  - originating server
  - user-specified time interval

Putting it into practice ...

Example:

White Consulting, Inc. has SLAs with all their customers, requiring them to guarantee mail delivery of all incoming messages within a specified number of minutes.

Using this PM Statistic type to gather verification data, White was able to prove compliance with the terms of their contracts.

PPM: Mail.Incoming.Delivery

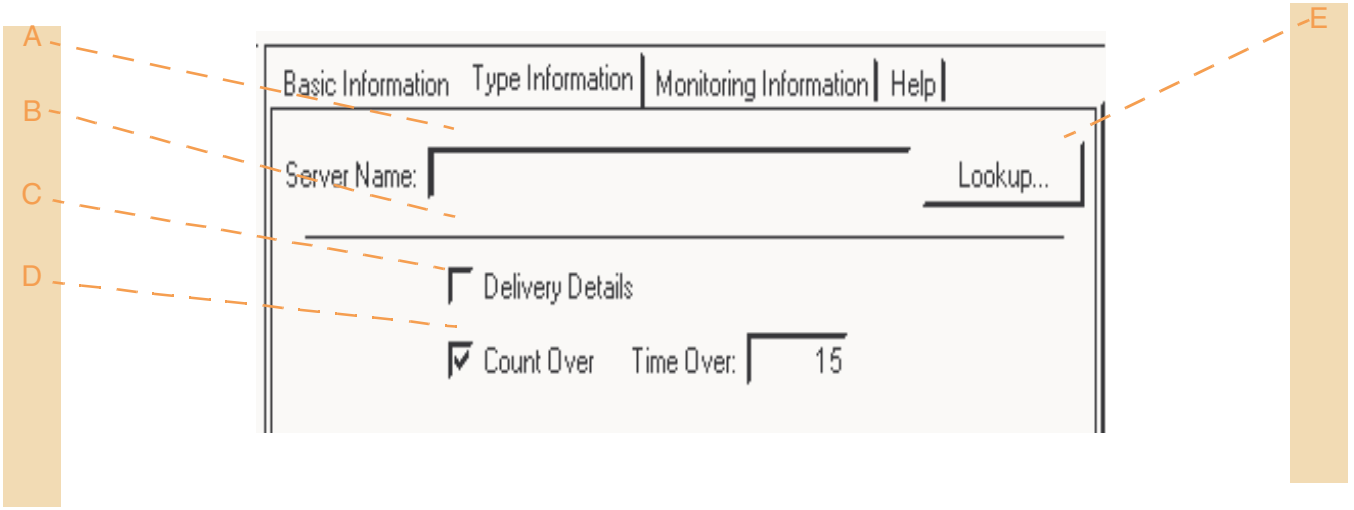

A: servers to monitor for this statistic
B: if selected, causes detailed delivery-time information to be generated
C: if selected, returns count of messages that took longer than the specified Time Over
D: delivery threshold (in minutes)
E: maximum allowable delay (in minutes) for mail delivery

PM Statistic Types: Mail.Outgoing.Attachment.Types

Basics ...

Reports:

- count and total overall size of attachments
- broken down by
  - file extension (type)

Thresholds:

The Size and Count thresholds (see fields B and C on the following page) apply to in-memory statistics only, and have no effect on what is reported to the IntelliWatch statistics repository **iwstats.nsf**.

Thresholds function cumulatively, and independently.

Example: (Count.threshold=100; Size.threshold=4MB)

If 99 messages are sent out, containing the targeted attachment type(s) and totaling 3.999KB, neither Count nor Size are available for in-memory monitoring. Values *are* written to **iwstats.nsf**, however, regardless of their count/size.

If 50 messages are sent out, containing the targeted attachments and totaling 5MB, the Size value *would* be available for monitoring, although the Count value would not.

Please bear in mind, however, that a Statistic Trigger checking for a value in units N/A (used for integers), won't fire based on this condition, since the monitored parameter is Count, not Size, whereas the Size value is the only one available in memory.

Putting it into practice ...

Example:

Server performance in your environment is suffering due to high volumes of mail messages with large attachments.

Use this statistic to discover just what types of attachments are having the most impact on server performance.

To obtain maximum benefit from this statistic, select it *both* for in-memory *Monitoring* and *Reporting.*

- *Monitoring* Mail.Attachment.Type with a Trigger alerts you to incidental issues involving large mail attachments.
- *Reports* on this statistic can provide data for a more detailed analysis of how attachment traffic is impacting server performance in your environment.

PPM: Mail.Outgoing.Attachment.Types

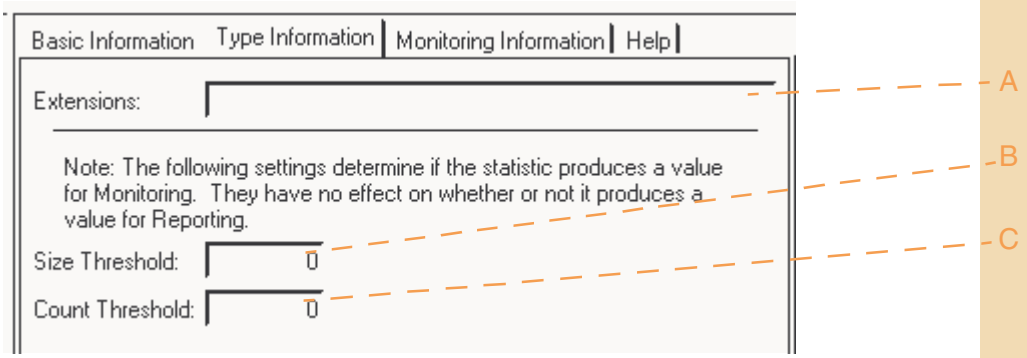

A: comma-delimited list of file extensions
B: size threshold of outgoing mail (applies to Triggers *only*)
C: count threshold of outgoing mail (applies to Triggers *only*)

PM Statistic Types: Mail.Outgoing.Server.Attachment.Percentage

Basics ...

Reports:

- percentage of outgoing mail with attachments
- broken down by:
  - destination server

Putting it into practice ...

Example:

Jones Electronics suspected that large attachments were adversely affecting mail server performance in their environment.

Using this PM Statistic type, they were able to ascertain:

- the percentage of outgoing mail with attachments
- where messages with attachments were being sent

Whenever an issue was detected, they used this statistic— in tandem with Mail.Attachment.Type—to isolate the source of the problem.

PPM: Mail.Outgoing.Server.Attachment.Percentage

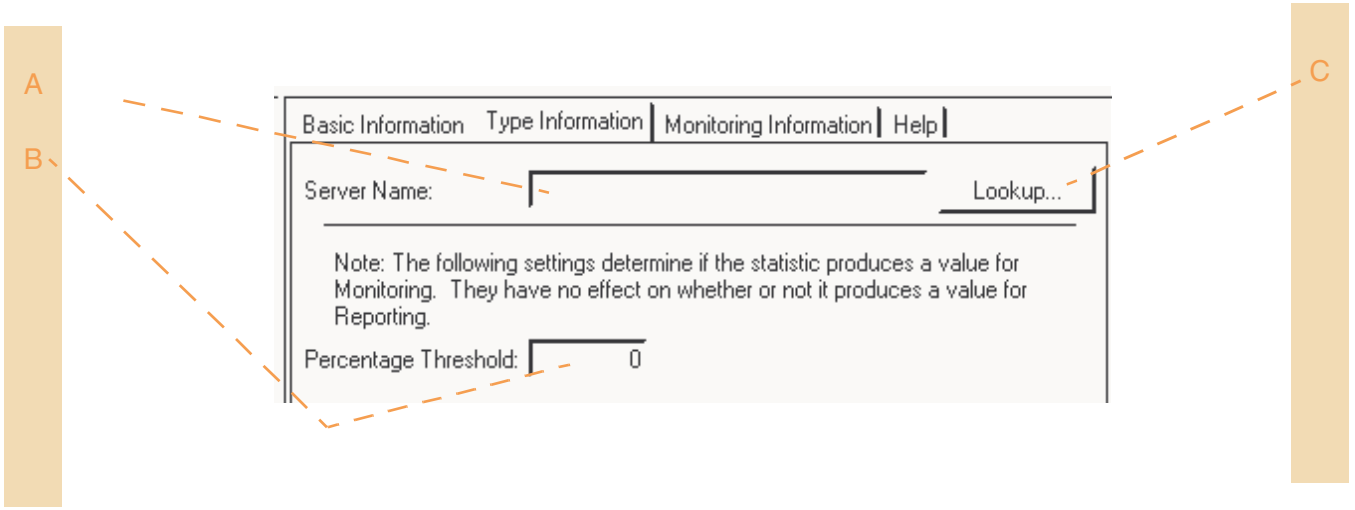

A: server(s) on which to create statistic (comma-delimited list)
B: threshold value (under which Triggers monitoring this statistic won't fire)
C: button to launch Select Server dialog

PM Statistic Types: Mail.Outgoing.Server.Volume

Basics ...

Reports:

- overall outgoing mail volume (in bytes)
- total number of messages sent
- broken down by:
  - destination server (next hop)

Thresholds:

The Size and Count thresholds (see fields B and C on the following page) apply to in-memory statistics only, and have no effect on what is reported to **iwstats.nsf**.

Thresholds function cumulatively, and independently.

Example: (Count.threshold=2000; Size.threshold=30MB)

If 1999 messages totaling 29.999KB are sent out, neither Size nor Count are available for in-memory monitoring. Values *are* written to **iwstats.nsf, however, regardless of their size/count**.

If 2001 messages totaling 5MB were sent out, the Count value *would* be available for monitoring, but not the Size value.

Statistic Triggers checking for a value in bytes would not fire on this condition; Statistic Triggers checking for units N/A (integers) would fire.

Putting it into practice ...

Example: 1:

Periodically, mail server load-distribution in your environment need to be reassessed.

This PM Statistic type generates data on outgoing mail volume that can assist you in making that assessment.

Example: 2:

Use this statistic type to track the volume of outgoing mail, and to identify potential abuse of your company's e-mail system.

The breakdown by destination server enables Admins to zero in on messages going to servers having nothing to do with company business.

PPM: Mail.Outgoing.Server.Volume

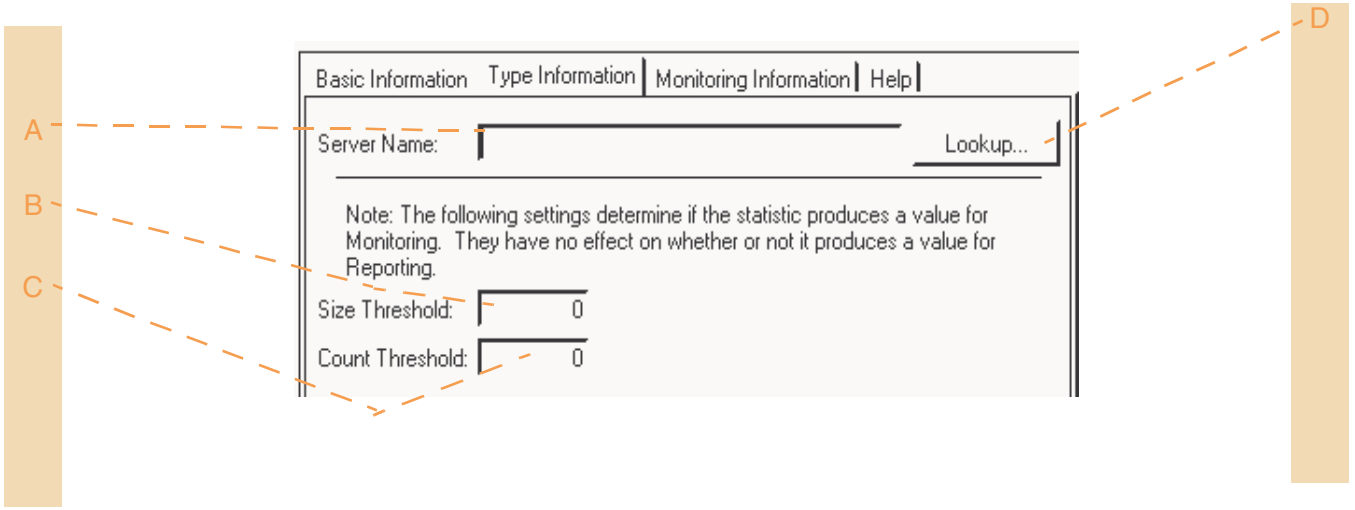

A: server(s) on which to create statistic (comma-delimited list)
B: size threshold of outgoing mail (applies to Triggers *only*)
C: count threshold of outgoing mail (applies to Triggers *only*)
D: button to launch Select Server dialog

PM Statistic Types: Mail.Size

Basics ...

Reports:

- User level reports
  - both sent and received mail, by individual User
- Server level reports
  - *outgoing mail only*, reported by destination server
- Maximum/Average Summary
  - reports a *total* of all selected users

To see mail activity broken down by individuals, choose the User-level report.

For all three types, reports:

- Count
- Max (maximum size, in bytes)
- Avg (average size, in bytes)

Control what is written to **iwstats.nsf** using the optional minimum reporting threshold. The threshold applies to a TOTAL of all messages sent (and received, when applicable). If the total number of bytes is not exceeded, the statistic will not be collected for the relevant user/server.

Putting it into practice ...

What is returned, by report option:

- User level reports

iwstats.STATISTIC_NAME.sent.USER.[max|avg|count]

iwstats.STATISTIC_NAME.sent.USER.[max|avg|count]

- Server level reports

iwstats.STATISTIC_NAME.outgoing.SERVER_NAME.[max|avg|count]

- Maximum/Average Summary

iwstats.STATISTIC_NAME.sent.[max|avg|count]

iwstats.STATISTIC_NAME.recv.[max|avg|count]

Example:

TriCounty Consulting, Ltd. specializes in the assessment of customers' Notes environments, particularly mail traffic.

One of their first steps on a new engagement is to set up Mail.Size statistics that enable them to quickly assess mail volumes, at the User and the Server levels.

A higher than expected average size of messages sent by certain Users suggested lines of inquiry, such as:

- Was average message size too high because Users were making inappropriate use of the *Reply with History* function (instead of simply *Reply*)?
- Was the maximum message size, while within the limits imposed by the customer's firewall, causing server performance issues?

Size: PM vs Notes:

Pinnacle reports the size of the message body, plus any attachments. Notes adds certain overhead items (making the sizes reported by Notes slightly higher).

PPM: Mail Size

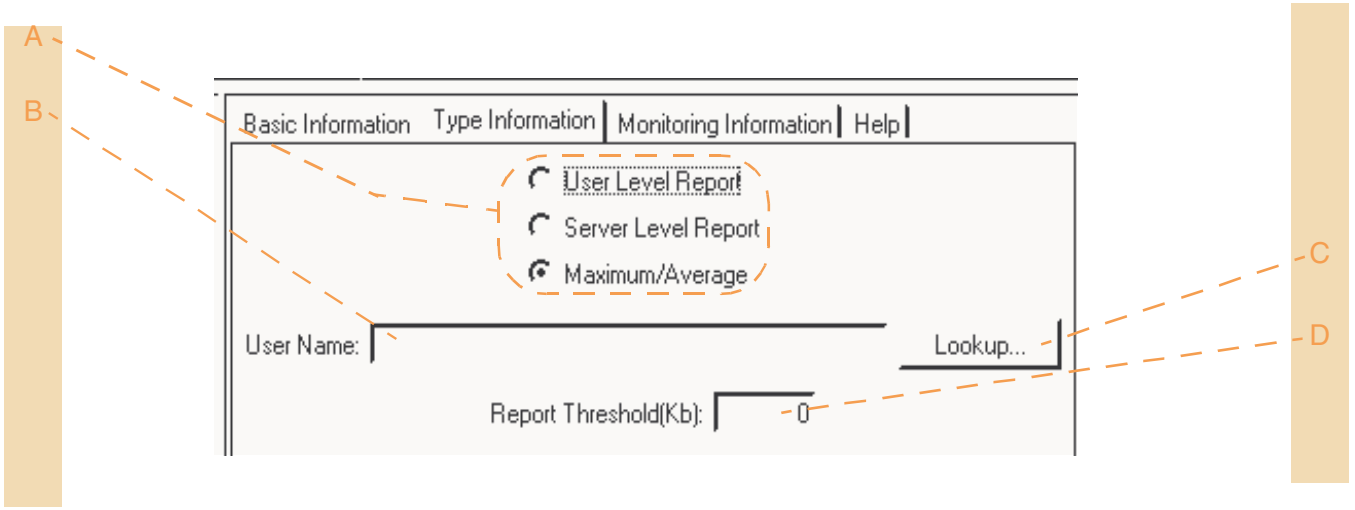

A: select report type
B: select users/servers to monitor, or summary information on the size of mail messages
C: Lookup button for users (pictured here)/servers on which to collect statistic
D: level below which the statistic are not collected for reporting purposes

PM Statistic Types: NT Performance Counter

Basics ...

- based on NT Performance Counters (NT only)
  - Any Counter maintained by NT can be used to create a custom statistic.
  - Returns average for the Counter over the period specified in the statistic's reset interval.
  - ***Note:*** If you define a statistic for a non-existent Counter, a value of zero is published.

Why is the Instance I want to monitor not listed?

One possibility is that the process is associated with a proprietary program. To add it to the list:

1. Enter the name of the process in the Unlisted field.
2. Save the statistic.

The next time you open that statistic, the Instance is included in the Instance list box, and the Unlisted field is blank.

Don't forget that a single program may have multiple processes associated with it. Think how may processes run under Notes!

Putting it into practice ...

Objects, Counters and Instances:

- Object: any system component with measurable properties
  - *physical component* (example: hard disk)
  - *logical component* (example: disk volume)
  - *software component* (example: process)
- Counter: measurable attribute of an object
  - *Processor* object includes Counters for *Interrupts/sec,* and *% User Time*, et al.
  - *Thread* object includes Counters for *% User Time,* and *Start Address*, et al.
  - *LogicalDisk* object includes *Disk Reads/sec* and *% Disk Read Time*, et al.
- Instance: occurrence of an object on a system

Example:

The *Processor* object represents a CPU (not a *Process*—which has its own object type). Therefore, a system with one CPU has:

- one *Processor* object (numbering starts at 0)
- several Counters associated with the object (such as *% Processor Time, DCP Rate, % Interrupt Time, % Privileged Time*, and so on*)*
- one *Instance* of the *Processor* object, namely the single CPU

(For additional information on Performance Counters, see http://windowsitlibrary.com/Content/169/01/5.html.)

PPM: NT Performance Counters

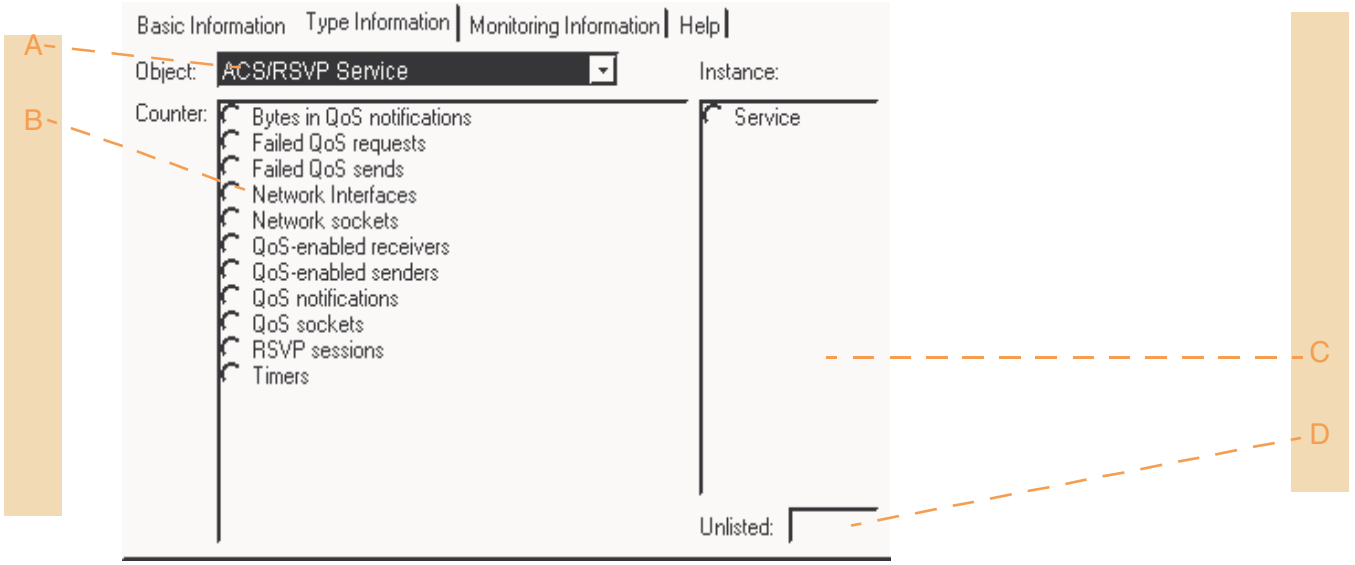

A: combo box for selection of object type (double-click icon for examples)
B: select radio button for desired Counter
C: instance list
D: type in an unlisted instance variable

PM Statistic Types: Replication

Basics ...

Reports (for push/pull):

- bytes sent/received
- docs transferred
- time taken
- number of successful replications per server
- number failed replications per server

Putting it into practice ...

The Replication statistic type requires no type-specific user input.

The statistic returns the same information (see previous page) for each server listed on the Monitoring Information tab (or for all known servers, if the field is left blank).

To configure this statistic type for Monitoring and/or Reporting, follow the same procedures as for all other IntelliWatch statistic types.

PM Statistic Types: Replication Delay

Basics ...

Reports:

- amount of time (in minutes) since replication last occurred
- broken down by:
  - server
  - database

Putting it into practice ...

Example:

Marketing Associates maintains a critical stock-tracking database. This Notes application provides information upon which significant financial decisions are made across the company. Replica copies of the database need to be kept in sync, especially when important decisions are about to be made.

Using the Replication.Delay statistic type simplifies Marketing Associates' replica-management procedures.

To obtain maximum benefit from this statistic, select it *both* for in-memory *Monitoring* and *Reporting*.

- *Monitoring* Replication with a Trigger alerts you when the (time) threshold you specify is exceeded.
- *Reports* on this statistic help you create a server-by-server *Replication Profile* for your environment.

PPM: Replication Delay

A

A: name of the database whose replication you want to check

PM Statistic Types: Server Event Count

Basics ...

Reports:

- a count of a specified Notes event

The statistic can be associated with a specific:

- add-in task
- user
- database
- URL (use when Server Event is Web Hit)

Multiple entries are comma-delimited, and wildcards are allowed (for examples, see *"Patterns:" on page 226*).

Associating Web Hits with a URL:

The value of a Pinnacle Web Hits statistic reflects the number of *successful* hits, not the number of attempts.

Web hits are not restricted by file type, but strings entered in the HTTP field must represent what appears in the browser's address field when accessing a document over the Web (or a portion thereof, when using wildcards).

Putting it into practice ...

When to specify optional parameters:

If all three optional parameters are left blank, all databases are checked, regardless of the User or Add-in task participating in the Server Event.

Using options wisely can help you zero in on which add-in task (or user) is interacting with which database, thus narrowing considerably the field of inquiry.

In the absence of clues, leaving all optional fields blank widens the field of inquiry and increases your chances of eventually tracking down the source of the issue.

Example::

World Widgets, Inc. recently had to request users of **Training.nsf** to access it on a different server, while maintenance was being done on the Applications Hub.

However, when employees were asked to resume accessing the database on the Applications Hub, many of them kept connecting to the Temporary Hub.

Just which employees were not being responsive to the request was established by configuring the statistic to register DbAccesses to **Training.nsf** on both the Applications Hub and the Temporary Hub.

Using this information, a second mailing was sent to only those employees who had not complied with the earlier request to resume accessing the Applications Hub.

PPM: Server Event Count

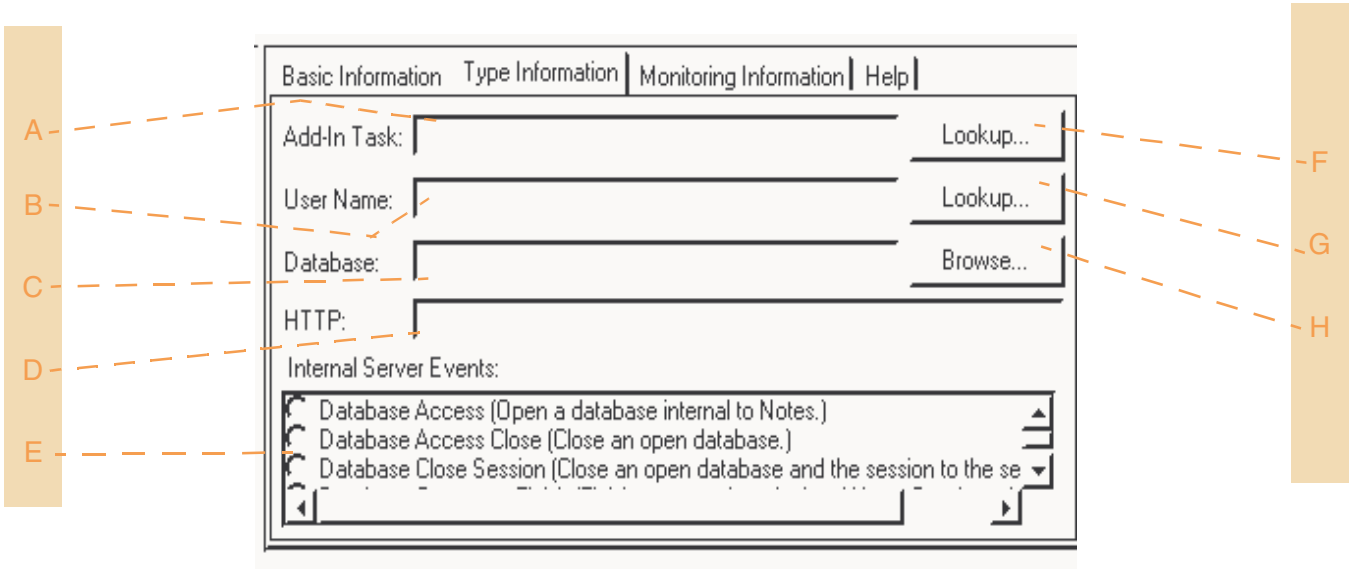

A: add-in task calling for server event (optional)
B: user name specifier (optional)
C: database name specifier (optional)
D: comma-delimited list of URLs (optional)
E: list of internal server events
F: button launches Select Add-in dialog
G: button launches Select User dialog
H: button launches Select Database dialog

PM Statistic Types: Summation

Basics ...

Reports:

- value of a single statistic
  - This allows you to give a Notes statistic a custom name, and to store it in the IntelliWatch statistic database.
- sum of values for two or more statistics
  - This allows you to combine the values of multiple statistics into a single Pinnacle statistic.

A list of commonly summed statistics is created by the Setup. This list is displayed on the statistic's Type Information tab.

If a statistic you want to use is not listed:

For information, see *"Adding unlisted statistics:" on page 218*, and *"Edit Notes Statistic List" on page 201*.

Putting it into practice ...

Example:

Black Moon Industries was in the process of reassessing the system resources needed to handle their Domino configuration.

To assess the peak physical memory requirements of individual server tasks, they created NT Performance Counter statistics for the Process object that tracked *Working Set Peak*.

By creating a PM Summation statistic, they were able to consolidate this data, and see at a glance the amount of physical memory needed by Domino.

PPM: Summation

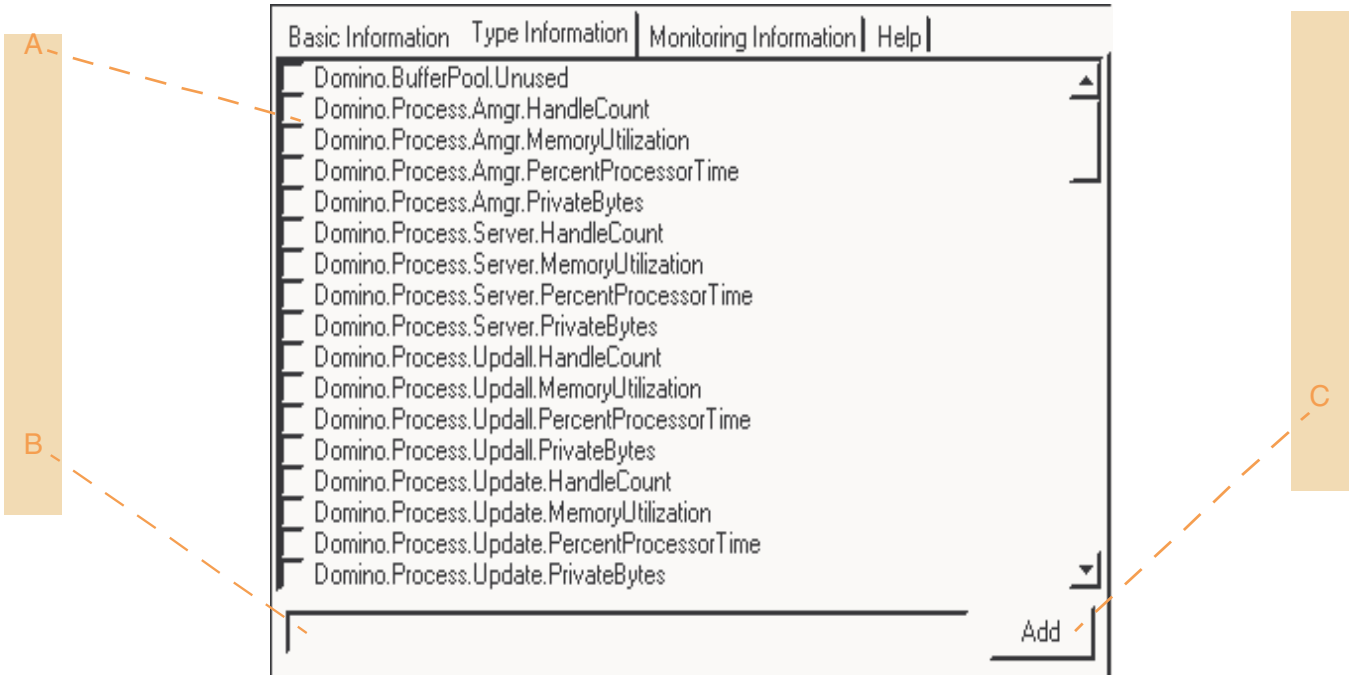

A: select statistics to be summed
B: text box for adding statistics not in list
C: click button to add statistic in text box to list

PM Statistic Types: View Performance

Basics ...

Offers three options:

- Server Summary
- Database Summary
- By View

For all three, four statistic values are returned:

- Avg, or Average (in milliseconds)
- Count
- Max, or Maximum (in milliseconds)
- PercentUnder[user-specified SLA threshold]=[%]

How views are incorporated into statistic names:

- Server Summary (no views are named)
  - Example: iwstats.[YourStat].Avg.
- Database Summary (no views are named)
  - Example: iwstats.[YourStat].Avg.
- By View
  - Example: iwstats.[YourStat].[ViewName].Avg.

Views in folders are displayed as: .[FolderName].[ViewName].

Putting it into practice ...

When interpreting values, bear in mind the interval applicable to the operation, reporting or monitoring.

Delimiting fields:

This statistic type allows you to filter instances on the basis of the User Name, Database or View.

Example:

Reporting of *iwstats.ViewPerf.DBSum* occurs hourly, but your Monitoring interval is 15 minutes. (For purposes of illustration, the SLA threshold is 1000[ms]).

At Point X in time:

The statistic in **iwstats.nsf**, reflecting an hour's worth of data, might have values such as the following:

- Avg=38
- Count=178
- Max=5000 (due to accessing a large database)
- PercentUnder1000=93

The in-memory statistic, representing only 15 minutes of data, might have values such as these:

- Avg=8
- Count=31
- Max=55
- PercentUnder1000=100.

PPM: View Performance

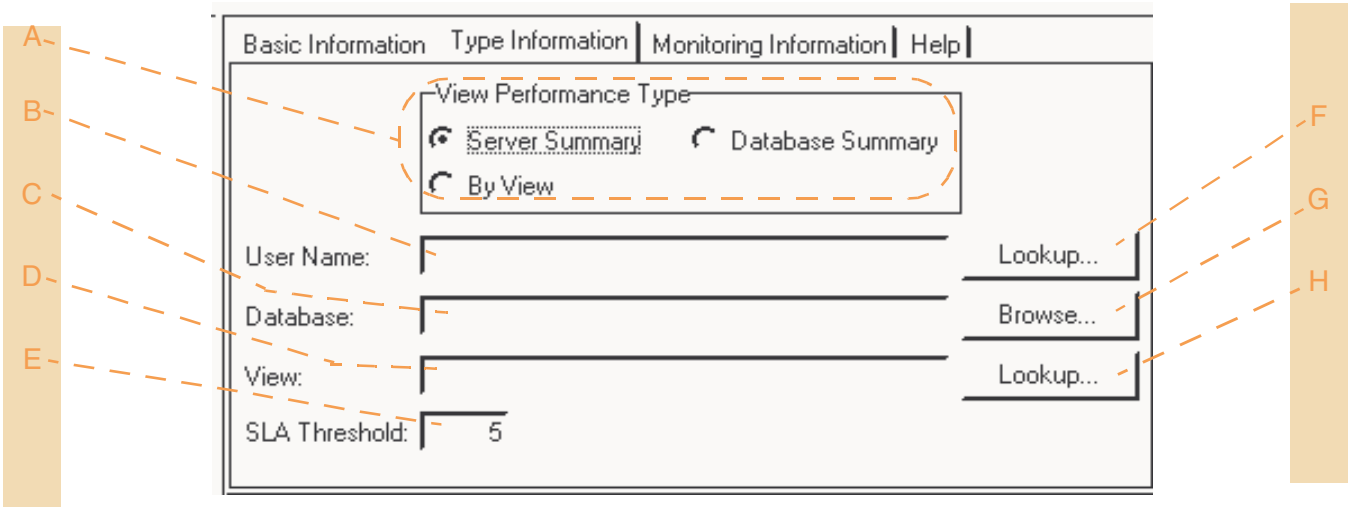

A: select View Performance type
B: User Name field (enabled only for Database Summary and By View types)
C: Database field (enabled only for Database Summary and By View types)
D: View field (enabled only for By View type)
E: applicable SLA threshold
F: button launches Select User dialog
G: button launches Select Database dialog
H: button launches Select View dialog

Chapter

Tracer

6

Use Tracer to diagnose bottlenecks in mail routing, database access, and replication, as well as for tracking intermittent load- and time-related Notes events.

Tracer records events in real time. Information is stored in a log for processing, then displayed at the Pinnacle Console. Saved Traces can also be opened at the Console for (re)analysis.

Chapter Contents

6.1.0.0 ARCHITECTURE

FIGURE 6-1: Tracer Architecture

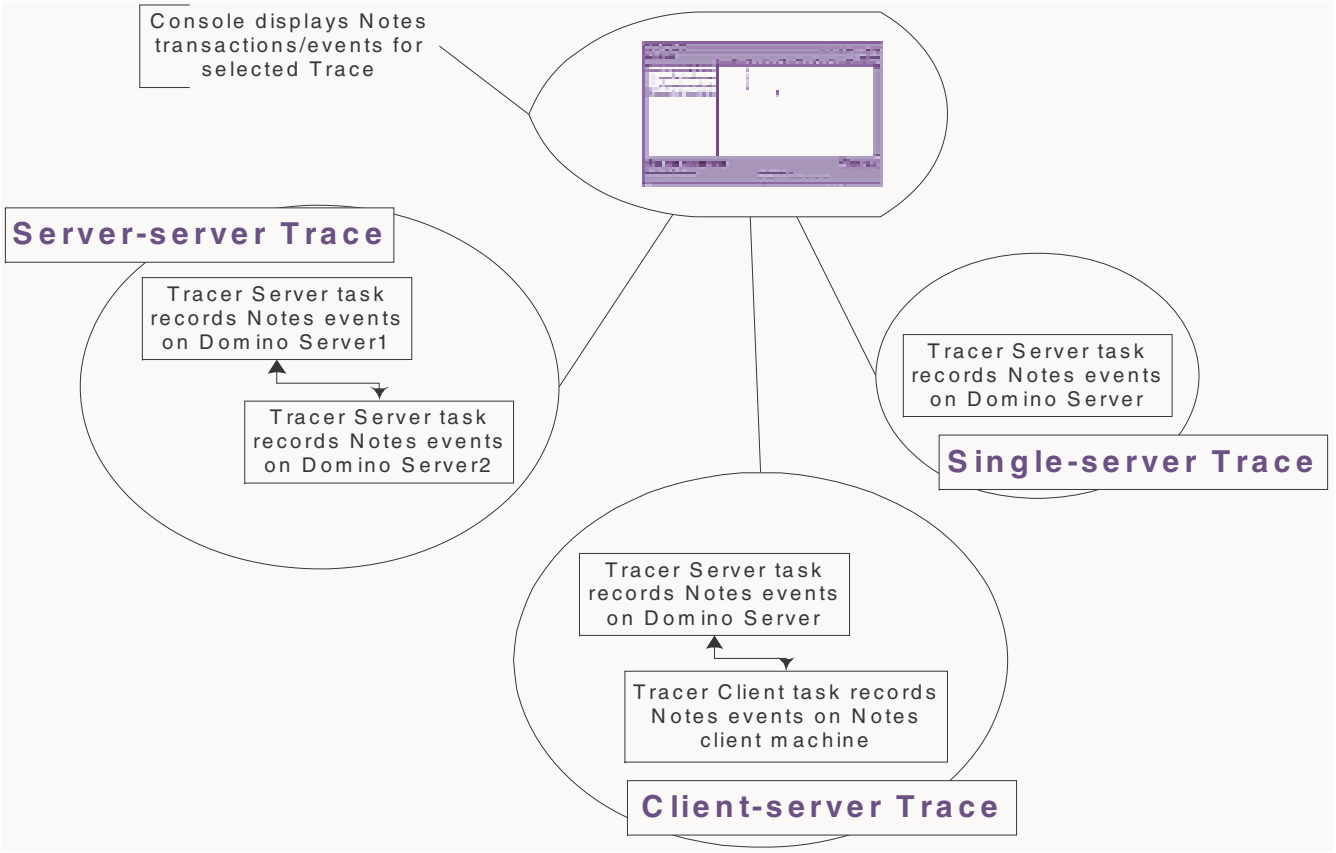

6.2.0.0 TRACER COMPONENTS

As shown in Figure 6-1, Tracer consists of three components: a server piece, a client piece, and the Tracer Solution at the Pinnacle Console.

- **Tracer Server**

  Running as a Notes add-in task, Tracer's server component must be loaded on each server that participates in a Trace.

- **Tracer Client**

  Required on any client you want to involve in a client–server Trace. This component records and transmits Notes client-side events to the Tracer Console.

  The Tracer client component runs on Windows 95/98 and NT platforms, and presupposes that a Notes client is installed on the machine.

  This is the only IntelliWatch component that you install on your Admin workstation (or other client machine you want to include in a Trace).

- **Tracer Console**

  Accessed at the IntelliWatch UI, the Tracer Console receives and graphically displays data. While the way data are displayed can be modified by means of the Tracer Console's views and filters; the data in the file remain unchanged.

6.3.0.0 TRACE TYPES AND THEIR USES

There are three types of Traces, each with a specific diagnostic purpose:

6.3.1.0 Single Server

- Useful for diagnosing overnight replication problems.

  Since many factors can contribute to replication delays, use Tracer to record (and time) all Notes events.

  A single-server Trace can be run overnight, but the large amount of data Tracer processes can cause delays. To minimize amount of data returned, limit the Trace to the most critical databases.

6.3.2.0 Client to Server

- Employ when users experience slow response times from Notes servers.

  A client-server Trace allows you to determine whether the delay is due to:

  – a slow network connection
  – a corrupted database (or view)
  – an unusually busy server.

6.3.3.0 Server to Server

- Use for diagnosing problems with server transactions, such as mail routing or replication.

  Poor response times between servers can be caused by the same issues that slow down client-to-server transactions (see above).

6.4.0.0 TRACER CONSOLE

The discussion of the Tracer Console is followed by an explanation of the dialogs used to run Traces

6.4.1.0 Running Traces

Configuring new Traces is a straightforward procedure.

6.4.1.1 Single server Trace

To diagnose local database performance on a Domino server, run a single server Trace.

Two parameters are involved in this type of Trace: the name of the server; the database(s) to be monitored, with the following options:

- specify a single database
- specify multiple databases (enter a comma-delimited list)
- enter *.* (or leave blank) for all databases

TO RUN A SINGLE SERVER TRACE:

1 Go to **File > New Trace** via the drop-down menus; alternatively, click on the toolbar icon.

This displays the **Start New Trace** configuration screen.

FIGURE 6-2: New Trace dialog

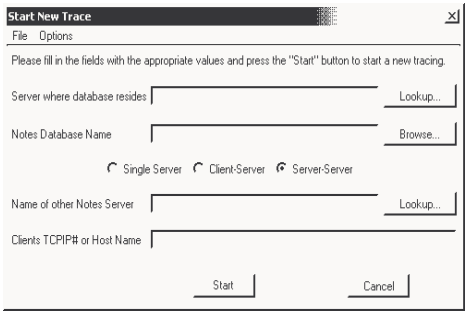

2 Enter or select the following information:

- Server where database resides (required)
- Notes database name (optional)

  Leave blank for all databases on the server.
- Server Only radio button (required)

3 Click the **Start** button.

This displays the **IntelliWatch Tracer** connection screen.

FIGURE 6-3: Tracer initialization dialog

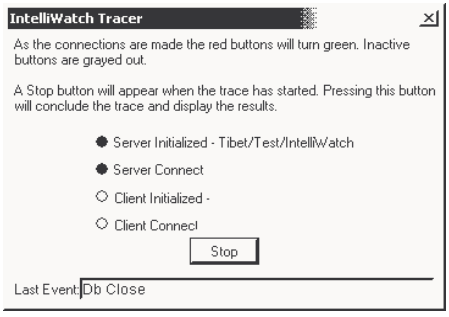

(While initialization is taking place, one or more initialization indicators remain red, and the button text reads Cancel. When initialization is complete, all lights are green, and the button text reads Stop.)

Note that only the top two initialization indicators come into play with a single server Trace. The lower two are greyed out.

4 Once initialization is complete, begin the operation where the delay is perceived (open the suspect database or view, for example). Alternatively, let Tracer run to monitor database activity as it occurs naturally.

5 Click **Stop** when the operation to be diagnosed has completed.

6 View the Trace.

The default is the Transaction view.

6.4.1.2 Client-server Trace

A client-server Trace is called for if performance problems are evident when a Notes client connects to a database application on a particular server.

There are four parameters, three of which are required:

- server name (required)
- database (optional)
  - specify a single database
  - specify multiple databases (enter a comma-delimited list)
  - enter *.* (or leave blank) for all databases
- user name (required)
- TCP/IP# or hostname of the client machine (required)

TO RUN A CLIENT-SERVER TRACE:

1 Go to **File > New Trace** via the drop-down menus; alternatively, click on the  toolbar icon.

This displays the **Start New Trace** configuration screen. (See Figure 6-2, above.)

2 Enter or select the following information:
- Server where database resides
- Notes database name
- Click Client-Server radio button
- User's hierarchical Notes name or client
- Client TCP/IP number or hostname

3 Click the **Start** button.

This displays the **IntelliWatch Tracer** connection screen.

initialization indicators remain red, and the button text reads Cancel. When initialization is complete, all lights are green, and the button text reads Stop.)

4 Once initialization is complete, begin the operation where the delay is perceived (open the suspect database or view, for example). Alternatively, let Tracer run to monitor database activity as it occurs naturally.

5 Click **Stop** when the operation to be diagnosed has completed.

6 View the Trace.

The default is the Transaction view.

6.4.1.3 Server-server Trace

A server-server Trace is called for if performance problems are evident with inter-server database application activity.

There are three parameters, two of which are required:

- name of first server (required)
- database (optional)
  - specify a single database
  - specify multiple databases (enter a comma-delimited list)
  - enter *.* (or leave blank) for all databases
- name of second server (required)

(Note that the field "Client's TCP/IP# or Host Name" is greyed out.)

TO RUN A SERVER-SERVER TRACE:

1 Go to **File > New Trace** via the drop-down menus.; alternatively, click on the  toolbar icon.

This displays the **Start New Trace** configuration screen (see Figure 6-2).

2 Enter or select the following information.
- server where database resides
- Notes database name

- server's hierarchical Notes name (use the pop-up dialog, or enter by hand)

3 Click the **Start** button.

This displays the **IntelliWatch Tracer** connection screen (see Figure 6-3).

initialization indicators remain red, and the button text reads Cancel. When initialization is complete, all lights are green, and the button text reads Stop.)

4 Once initialization is complete, begin the operation where the delay is perceived (open the suspect database or view, for example). Alternatively, let Tracer run to monitor database activity as it occurs naturally.

5 Click **Stop** when the operation to be diagnosed has completed.

6 View the Trace.

The default is the Transaction view.

6.4.1.4 Saving Traces

Traces are often attempts at diagnosis, that can be discarded when a particular line of inquiry has proven unfruitful.

When diagnosis is not possible at the time a Trace is run, or if you want to save Traces for later review and analysis, use the following procedure.

TO SAVE TRACES:

1 Go to **File > Save** via the drop-down menus; alternatively, click on the  toolbar icon.

This displays the Save File dialog.

2 Enter a filename and location.

The default on NT 4.0 is:

WINNT/Profiles/Administrator.000/IntelliWatch/data

The default on Windows 2000 is:

Documents and Settings/administrator.INTELLIWATCH/IntelliWatch/data

3 Click OK to save the Trace.

Compose a filename that conveys the Trace's purpose to anyone who opens it.

If a filename is not entered, Tracer supplies a numerical string derived from the Trace's timestamp (979232136019.trc, for example).

6.5.0.0 FILTERS

Once a Trace has been run, the Console offers filtering at two levels, for:

- Database
- Notes Name
- Process Name

6.5.0.1 Level 1: Which events display

Single Server Traces are often run with the Database field left blank (or wildcards are used), so that Events associated with any database or user are recorded.

Once the Trace has been generated, first expand all Events/Transactions and search for particularly long event-times. If you determine that an issue is limited to a given database, use Level 1 filters to display only Events involving that database.

Next, survey all Notes Names (IDs) and Processes that have interfaced with the database in question, to see if the issue can be further narrowed down. Use the Notes

Name and Process filters progressively to limit what is displayed to only those Names/Processes that appear to be the source of the issue.

Using the combo boxes (see letters A-C at *"Filters" on page 271*), select those databases, Notes Names (read IDs) and processes whose initiated events you want to view.

6.5.0.2 Level 2: Amount of detail

This level of filtering governs the amount of detail in displayed events.

If all three options are selected (see the checkboxes at letter E of *"Filters" on page 271*), each recorded event includes:

- name of database involved
- Notes Name (ID) initiating the event
- process initiating the event

The following DbAccess Event illustrates what is displayed.:

DbAccess (Tibet/Test/IntelliWatch) (events4.ntf) (Fixup)

In this instance, the Notes Fixup task accessed **events4.ntf** on the Domino server Tibet/Test/IntelliWatch.

There is redundancy at the Transaction level, however, when the Process Name option is selected (see letter E, Show Additional Information, at "Filters" on page 271). The Process Name is displayed both before the word 'Event', and in parentheses at the end of the Transaction.

Deselecting the Process Name option causes it to be displayed only before the word Event (and not appended in parentheses).

However, thus configured, the Tracer Console no longer displays the Process Name if and when you switch to the Events View (see letter G at "Tracer Console" on page 261). To display the Process Name in the Events View, select it at letter E of "Filters" on page 271).

6.6.0.0 INTERPRETING TRACES

Interpreting Traces is more complex than running them. Just *how* complex is a function of how selective you are able to be in isolating a particular database as the source of an issue.

Since running and interpreting Traces can take a significant amount of time, we suggest you spend a few minutes before running a Trace to answer the following questions:

- Is the issue chronic or intermittent?
- Is there a logical explanation for the behavior you're witnessing?
  - Are all users experiencing slow connection times to the same server (or database), for example, or only certain users?

If only certain users are involved, what commonalities can you discern between them?

- Have design changes been made recently that have had a noticeable impact on database performance?

- If the issue is slow response time from a database in another domain, are *all* databases on that same server exhibiting similarly long response times?

Analyze Traces by making full use of the Filtering options discussed above.

6.7.0.0 FAQs

Q: Why do client-server Traces sometimes show server events preceding the first client event?

A: Possible explanations:

- Some events on the server are caused by Notes client events that are not trappable (Tracer is therefore not aware of them).

  This may relate to the manner in which a Notes client uses cache.

- Server events related to client activity may already be in progress when the Trace is started.

To limit unwanted events from appearing in a Trace:

- Use the Zoom-in feature to trim off the extra server events.
- Start the Trace just after launching the notes client you wish to Trace (do NOT open any databases before starting the Trace).

Tracer Console

Basics ...

Displays:

- Notes Events/Transactions pane
  - If the Events radio button is selected at the lower right of the Console (see letter G on the following page), Notes Events are displayed individually.
  - If the Transactions radio button is selected at the lower right of the Console, Events are grouped by Transactions.
- Timeline
- Name of server(s) participating in Trace

  See letter D on the following page.
- Database whose events are being recorded

  Includes time and date when Trace was initialized. See letter F on the following page.

How do Events differ from Transactions?

Events represent the 'steps' in a Transaction.

Assume all we want to do is open and close **ABC.nsf**: two Events, *DbOpen(ABC.nsf)* and *DbClose(ABC.nsf)*, make up one Transaction, *AccessDB*.

Putting it into practice ...

Toolbar:

In addition to the standard New, Open and Save icons, the Tracer toolbar also includes 'zoom' icons to allow you to reduce/expand the range of the Trace that displays in the Timeline pane.

TO ZOOM IN:

1. Click on the  icon.
2. Position the cursor over the left-most point on the Timeline to be included in the partial Trace.

   The cursor becomes a crosshatch.
3. Depress and hold down either the left or the right mouse button while moving the cursor from left to right—from right to left does not work— until you reach the right-hand extremity of the partial Trace.
4. Release the mouse button.

   The right-hand pane adjusts to display only the events selected.
   Multiple 'zooms in' are possible, if you need finer granularity.

TO ZOOM OUT:

1. Click on the  icon.

   This button is disabled until one *zoom in* has been performed. Once the icon is enabled, click once for each *zoom in* performed on the Timeline. Once back at the Trace's original Timeline, the icon is again disabled.

Color-coding:

- Client events are blue
- Server events are green
- Transactions are gray
- Triggered events (exceeded thresholds) are red

Tracer: Tracer Console

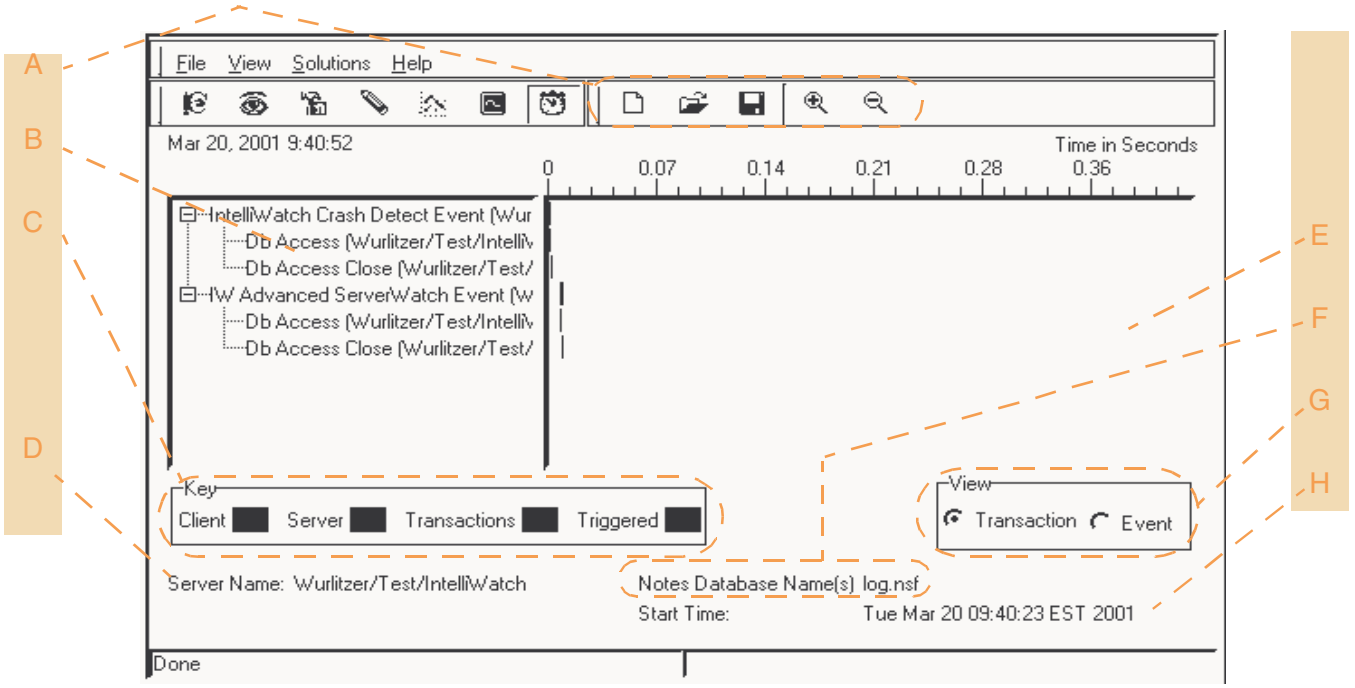

A: Tracer toolbar
B: list of events/transactions
C: color key for Trace analysis
D: name of server(s) involved in Trace (here on a single-server)
E: pane for time-line display of events (see scale at top of pane)
F: name of database(s) being traced
G: radio buttons to select Transaction or Event view (affects display at letter B)
H: start time of Trace

Single Server Trace dialog

Basics ...

Records:

- Notes Events between a
  - single Domino server, and
  - one or more databases *on that server*

Wildcards are supported. Leave blank for all databases on the server.

Putting it into practice ...

Fields requiring a value:

Only the top two fields require (or indeed accept) a value at this dialog, even though the background color of the "Name of other Notes server" field (white) looks as if it, too, should be filled in.

This behavior is due to a quirk of Sun Microsystem's implementation of Java, which causes the arrow of the combo box to gray out, but not the field itself. Functionality is not affected, and the field should be ignored.

Tracer: Single Server Trace

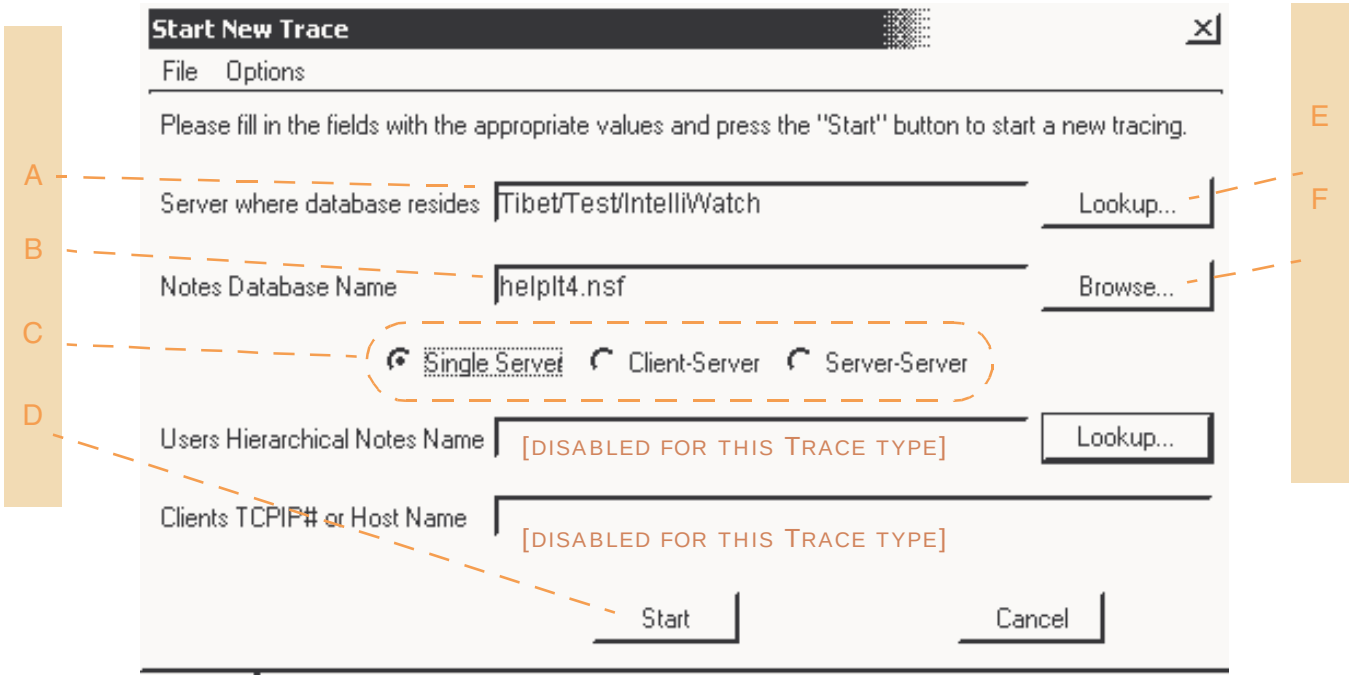

A: server involved in Trace
B: database(s) involved in Trace
C: type of Trace
D: start Trace button
E: button launches Select Server dialog (see "Selection dialogs," on page 268)
F: button launches Select Database dialog (see "Selection dialogs," on page 268)

Client-server Trace dialog

Basics ...

Records:

- Notes Events between a
  - a single client, and
  - one or more databases on a Domino server

Wildcards are supported. Leave blank for all databases on the server.

Putting it into practice ...

User information:

Three of the fields to be filled in on this dialog are self-explanatory:

- name of server where database resides,
- name of database(s)
- hierarchical name of user

The fourth field is somewhat less obvious:

- client's TCPIP# or Host Name

Before client-server Traces can be executed, obtain one or the other from your IT team.

Tracer: Client-server Trace

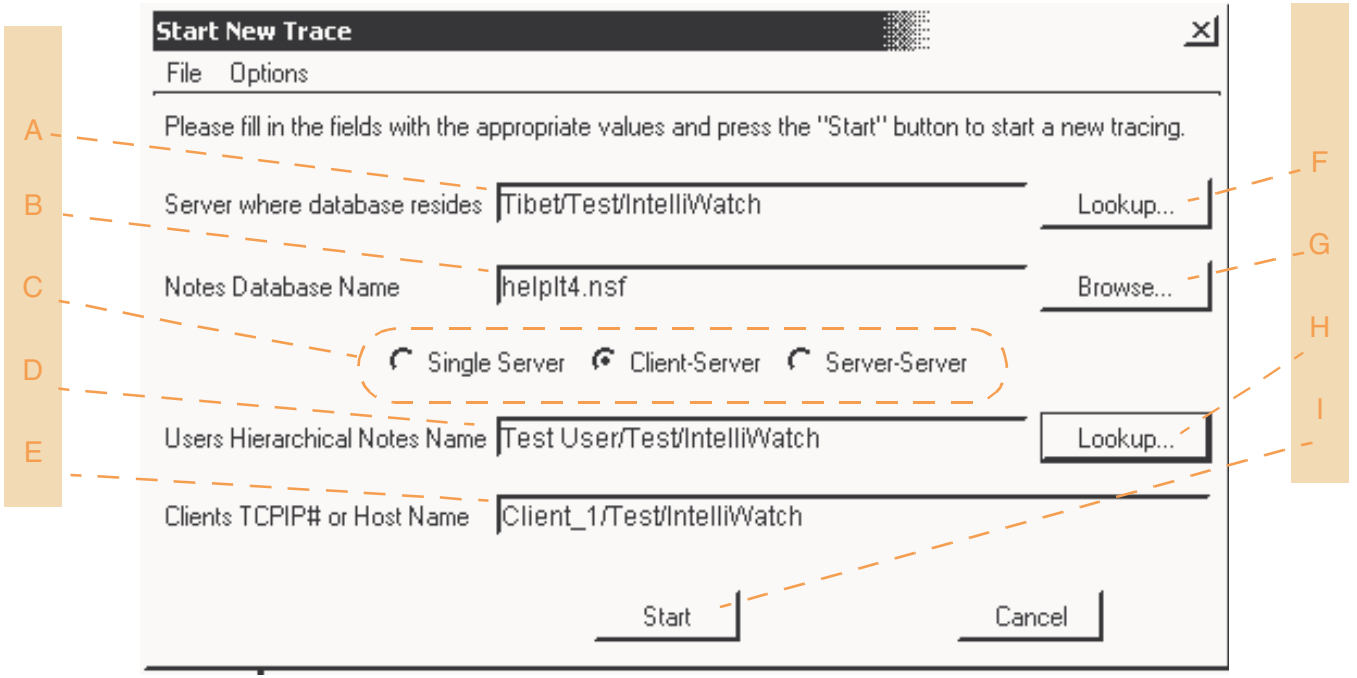

A: server involved in Trace
B: database(s) involved in Trace
C: type of Trace
D: Notes name of User involved in Trace
E: TCPIP# or Host name of client machine
F: button launches Select Server dialog (see "Selection dialogs," on page 268)
G: button launches Select Database dialog (see "Selection dialogs," on page 268)
H: button launches Select User dialog (see "Selection dialogs," on page 268)
I: start Trace button

Server-server Trace dialog

Basics ...

Records:

- Notes Events between a
  - a single Domino server, and
  - one or more databases on a second server

Wildcards are supported. Leave blank for all databases on the server.

Putting it into practice ...

Remember what is—and isn't—being recorded:

Only those events initiated

- by the server in field D (on the following page)
- on the database(s) in field B

are recorded.

Events initiated on the(se) database(s) by other clients or servers are not recorded (nor are events recorded involving unlisted databases on these two servers).

Tracer: Server-server Trace

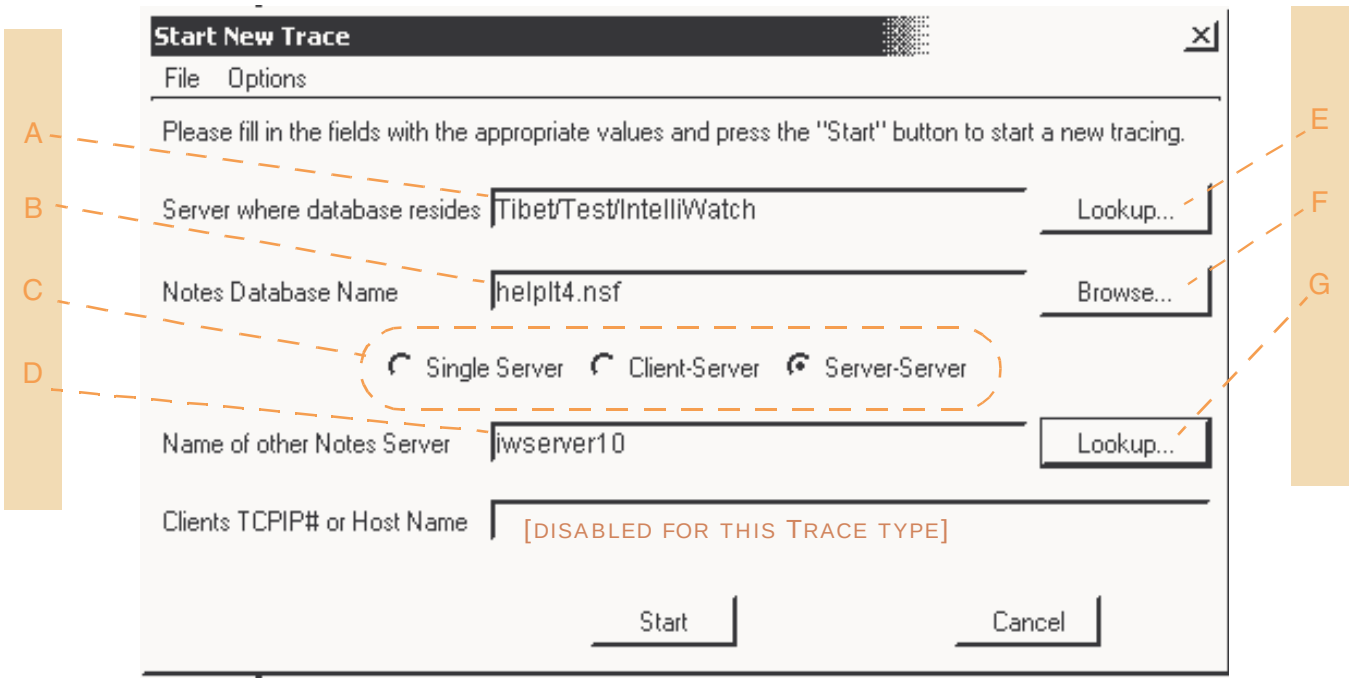

A: first server involved in Trace
B: database whose events are to be recorded
C: radio buttons to select type of Trace
D: second server involved in Trace
E: button launches Select Server dialog (see "Selection dialogs," on page 268)
F: button launches Select Database dialog (see "Selection dialogs," on page 268)
G: button launches Select Server dialog (see "Selection dialogs," on page 268)

Selection dialogs

Basics ...

Use to select

- Servers
- Databases
- Users

The Select Server dialog is pictured on the following page, but all three dialogs function similarly.

See 'Putting it into practice' for usage instructions.

Putting it into practice ...

Add/Remove buttons:

To move a server/database/user from the Choices to the Selected list box, click on it and then on the Add button. Alternatively, double-click your selection.

To remove a server/database/user from the Selected list box, click on it and then on the Remove button. Alternatively, double-click your selection.

User-specified:

To select a server/database/user not displayed under Choices, click the User-specified button. This brings up a dialog where you can enter the unlisted selection. Complete the selection process by clicking OK. (Click Cancel to abort the user-specified addition.)

A user-specified selection causes an error dialog to appear at the initialization stage of the Trace if the server, database or user cannot be found.

Similarly, even if the server, database or user exists, if the Console is unable to connect to it, an error results.

Remember ...

All servers involved in a Trace must be running the Tracer server task, or the Tracer client executable, for the operation to succeed.

Tracer: Selection

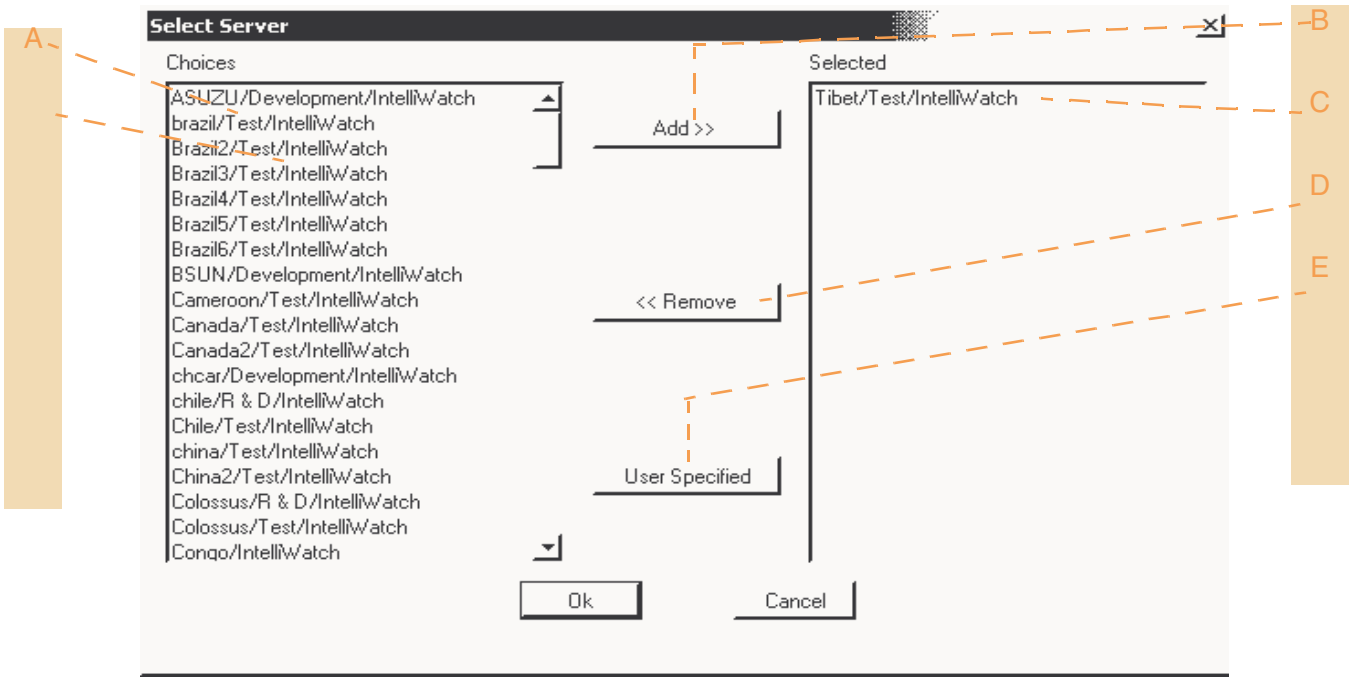

A: list of known servers
B: button to move server from Choices to Selected list box
C: selected server
D: button to remove server from Selected list box
E: button launches dialog for adding unknown server

Filter dialog

Basics ...

Use to refine Traces at the level of:

- Databases
- Notes Names (IDs)
- Processes

For usage details, see *"Filters" on page 256*.

Putting it into practice ...

Filtering levels:

Two levels of filtering are available:

- which Events are displayed
- amount of detail per Event

For usage details, see *"Filters" on page 256*.

Type:

The checkboxes which indicate 'type' (see letter D on the following page) are not editable; what is checked (though grayed out) reflects the elements of the Trace.

Tracer: Filters

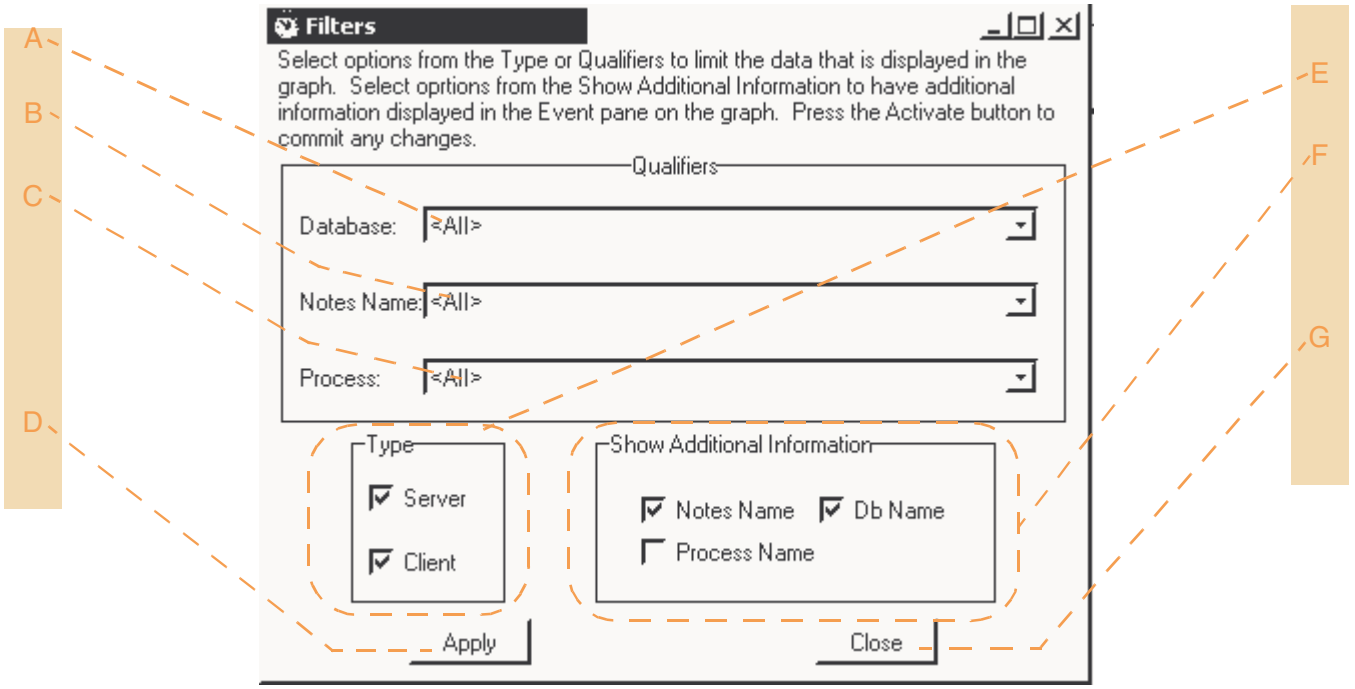

A: click combo box arrow to limit to Events associated with a single database, or leave at ALL (default)

B: click combo box arrow to limit to Events associated with a single Notes Name (ID), or leave at ALL (default)

C: click combo box arrow to limit to Events associated with a single process, or leave at ALL (default)

D: click Apply to filter Trace

E: non-editable indicator of elements participating in Trace

F: checkboxes to select/deselect additional thread information

G: click Close to return to the Tracer Console

Chapter

7

IntelliWatch Messaging Center Gateway

IntelliWatch Messaging Center's sole function is to relay messages from IntelliWatch components.

The application's simplicity makes for amazing flexibility. IwMCG receives a message over TCP/IP, and, based on internal variables, determines both message type and content. The Message Center Gateway then relays the notification in a form appropriate to the recipient/application.

Chapter Contents

7.1.0.0 Architecture

FIGURE 7-1: Architecture of IntelliWatch Messaging Center

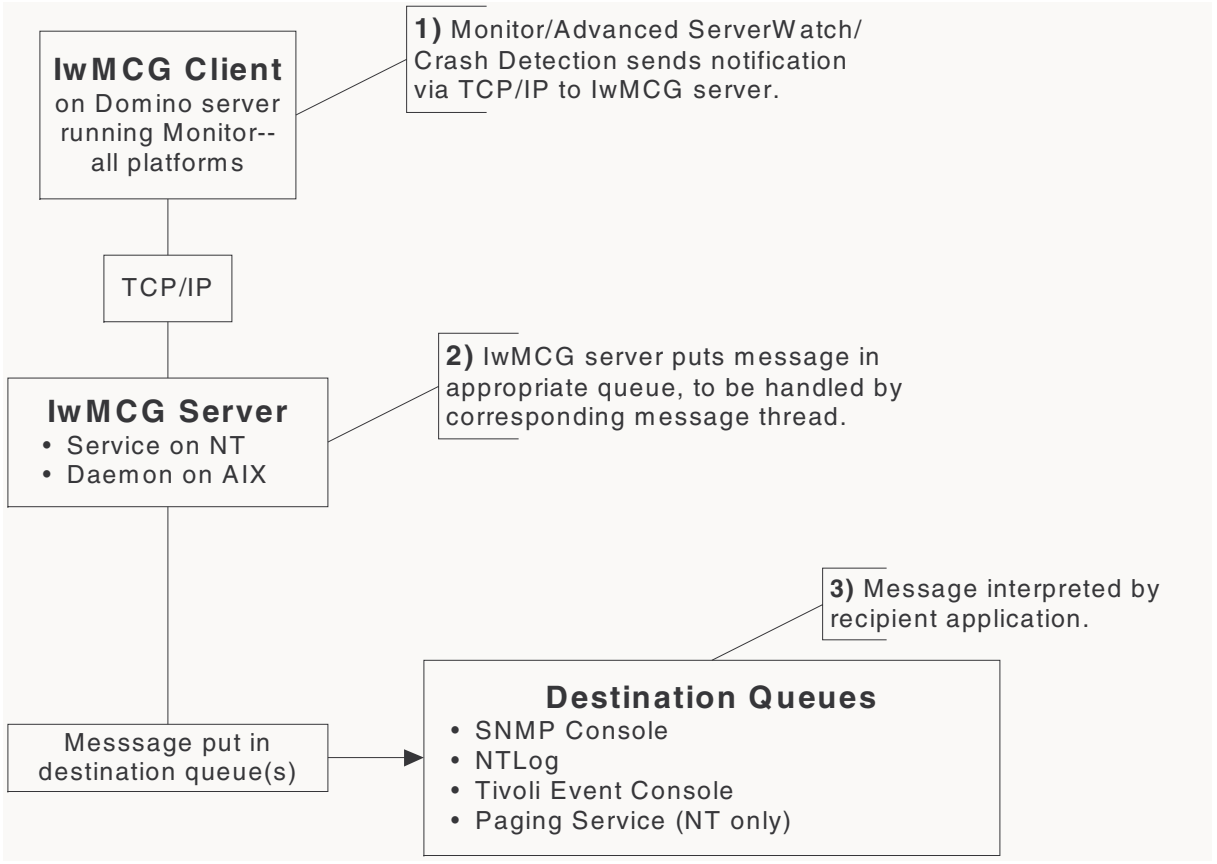

7.2.0.0 SENDING MESSAGES TO THE IWMCG

Messages can be sent to the Message Center Gateway by:

- command-line utilities
- Commands launched by IntelliWatch Monitor Triggers

7.2.1.0 IWEVENT

Three of the four message types handled by the Message Center Gateway use **iwevent**, a command-line utility that is new in IntelliWatch. Use it to send:

- SNMP traps (msg type=SNMP)
- NT Log Events (msg type=NTLog)
- Tivoli TEC Events (msg type=TEC)

To send a message, use this form:

```
iwevent [msg type] [other arguments]
```

To display type-specific arguments, type "iwevent [msg type]" at a command prompt, and press enter.

*Alphanumeric pages are sent, as in previous versions of IntelliWatch, by invoking **iwpage**.*

The former IntelliWatch Paging Server has been subsumed under the Message Center Gateway, however.

Paging is only supported when the Gateway is running on NT.

7.2.1.1 Tivoli TEC Events

No default TEC Event Command is included in the Command database (iwcmd.nsf). To create a TEC Event Command, use either the Pinnacle Console (see *3.4.1.2 on page 58*) or a Notes client (see *10.2.2.2 on page 330*).

7.2.1.2 Using %1 as a argument in a batch file

%1 can be used in connection with **iwevent** to pass the string corresponding to the error condition that caused a program to be launched.

Let's say you want Advanced ServerWatch to send a WARNING to the Tivoli Event Console when one or more of a certain group of servers are deemed NOT RESPONDING.

1. Create an Action Profile for those servers in Advanced ServerWatch (see *"Action Profiles" on page 179*).
2. Select Program as one of the options, and fill in the location of the following simple batch file:

```
iwevent TEC /C:WARNING /M:%1
exit
```

If Advanced ServerWatch determines that a server covered by the Action Profile is in a state that should cause that profile to be executed, the batch file is launched. A TEC Event (WARNING) with the relevant message is then sent, via the Message Center Gateway, to the Tivoli system listed in the *ServerLocation* parameter in the tecad_iw.conf file on the system where the batch file was launched.

At the Tivoli Console, %1 is replaced by the text of the message generated by Advanced ServerWatch.

%1 can also be used in Commands to be launched by a Monitor Trigger.

7.2.1.3 Alphanumeric pages

To send alphanumeric pages via the Message Center Gateway, enter:

```
iwpage [other arguments]
```

Remember ...

To obtain arguments for the "iwevent" utility, enter "iwevent [type]"--without additional arguments--at a command prompt.

To obtain arguments for sending alphanumeric pages, enter "iwpage"--without additional arguments--at a command prompt.

7.2.2.0 IntelliWatch Monitor Triggers

All message types handled by the Message Center Gatewaycan be sent by IntelliWatch Triggers.

7.2.2.1 Default Commands and the Message Center Gateway

The Commands invoked by IntelliWatch Monitor Triggers are contained in **iwcmd.nsf**.

Although this database includes Commands for all four message types handled by the Message Center Gateway (see Destination Queues in *Figure 7-1*, above), only TEC Event Commands and alphanumeric pages are *always* handled by the Message Center Gateway.

Messages sent to the NT Event Log, or as SNMP traps, can be sent either via the relevant local functionality or via the Message Center Gateway.

7.2.2.2 NT Event Log Commands

When Triggers launch the default NT Event Log Command, for example, **iwntlog** is invoked, and Events are sent to the *local* Event Log on the system where that Command is executed.

NT Log Events sent via the Message Center Gateway, on the other hand, are deposited in the Event Log *of the Gateway system itself*. If this is the desired destination for the NT Events, create a custom NT Event Log Command as follows:

Name: [short, meaningful name]
Program: iwevent
Arguments: NTLog [other arguments]

7.2.2.3 SNMP-trap Commands

Default SNMP-trap Commands use local SNMP functionality (running as a Service on NT, or as a daemon on AIX). The default Command invokes the executable **iwsdtrap**.

Create a custom Command to send SNMP traps via the Message Center Gateway as follows:

Name: [short, meaningful name]
Program: iwevent
Arguments: SNMP [other arguments]

While the Message Center Gateway can be used to send IntelliWatch SNMP traps on all platforms, it is most relevant for Solaris systems, since this is the only means of sending IntelliWatch traps on that OS (see 7.3.3.0).

7.3.0.0 PLATFORM-SPECIFIC CONSIDERATIONS

7.3.1.0 General

- IwMCG client piece runs on all platforms on which Monitor is supported.
- IwMCG server runs only on Win NT/2000 and AIX.
- TCP/IP must be available on both client and server.

7.3.2.0 AIX and alphanumeric paging

- IwMCG on AIX *does not* support alphanumeric paging.

  However, AIX servers running IntelliWatch Monitor can still send pages through a Messaging Center (or IntelliWatch Paging Server) running on NT.

7.3.3.0 Solaris and SNMP

- Solaris systems running Monitor are now able to send SNMP traps, by taking the following steps:

1 Install the Messaging Center server on an NT or AIX machine.

2 Create a new send-trap Command in the IntelliWatch Commands database (**iwcmd.nsf**).

3 Edit Triggers that you want to send traps by selecting the newly created Command.

7.4.0.0 SNMP MISC.

7.4.1.0 Native vs MCG

AIX and NT/Windows systems can send traps either natively (assuming the requisite system-level configuration is in place) or via the MCG.

Native traps use the iwsdtrap utility (as do those sent by Advanced ServerWatch, for example), whereas traps can be sent via the MCG using either the **iwsnmpevt** utility, or **iwevent** with **SNMP** as the first argument.

Solaris systems must use the MCG to send traps, as IntelliWatch does not support native SNMP on that platform (see also *7.3.3.0*, above).

7.4.2.0 SNMP Sets and gets

IntelliWatch SNMP traps—whether sent via the MCG or using native SNMP where available—do NOT support sets and gets.

7.5.0.0 MCG-ONLY INSTALLATION

7.5.1.0 Authorization code needed?

When the Message Center Gateway is the only IntelliWatch component installed on a given system, no authorization code is required.

7.5.2.0 Is Notes/Domino required?

No.

The Message Center Gateway can be installed on a supported OS (Windows or AIX) where neither a Notes client nor a Domino server is installed.

A drawback of this approach, however, is that the MCG cannot then be configured using the IntelliWatch Parameter Configuration utility (which requires a Domino server—NOT a Notes client—see also section *11.2.8.0 on page 353*).

7.5.3.0 Multiple versions unsupported

It is NOT supported to run a version of the Message Center Gateway that differs from the version of IntelliWatch running on monitored servers—for example, running 27.28 on monitored servers, but the 27.36 on the MCG itself..

Chapter

Crash Detection

8

IntelliWatch Crash Detection monitors server availability locally, complementing the remote connectivity monitoring provided by Advanced ServerWatch.

The simple-but-elegant local monitoring functionality is based on the ability of Crash Detection to open a local database (the default is **Log.nsf**). If the database can be opened, the server is considered to be running; if the database cannot be opened, Crash Detection follows a series of steps (the exact sequence of

Chapter Contents

which depends on various configuration options).

8.1.0.0 OVERVIEW

Crash Detection's overall paradigm is based on the simple question: "Can [database].nsf be opened?" If the answer is 'Yes', the server is available; if 'No', the server may have (has) crashed.

Numerous parameters can be customized, allowing you to configure Crash Detection to the needs of your environment. Options include:

- *Wait Times*
  - Two-stage *Wait Times* allow Crash Detection to adjust the number of seconds available to open the target database (before a crash is signalled). (For details, see Table 8-1 on page 287.)
- *Idle Time*
  - Disables Crash Detection during certain times of day.
- Notification methods
  - send page
  - send SNMP trap
  - run custom program(s)

(*.bat files are currently the ONLY supported type of custom program)

- action sequences
  - Notify only
  - Recycle, but do not reboot system if unsuccessful
  - Recycle, and reboot system if necessary
  - Reboot system

8.1.1.0 Database monitored

Most commonly, **log.nsf** is the monitored database.

While any Notes database on the local server can be monitored, a database that is automatically configuration-checked by Notes is preferable. Otherwise, a database that cannot be opened due to corruption, for example, would lead Crash Detection to assume the server had crashed, leading it to take unwarranted actions.

8.1.2.0 If database available

If Crash Detection is able to open the target database before the end of the *Initial Wait Time* (for a description, see *8.1.3.2*, below), no actions are taken, and the program awaits the start of the next checking cycle.

8.1.3.0 If database *un*available...

For step-by-step Crash Detection program flow, see *Figure 8-1*.

In global terms, Crash Detection works through the following steps when the monitored database application is

unavailable (in the following list *only*, consider *NOTIFY* to include all configured notification/launch program actions):

- check for enabled *Idle Time* (if any)
  - If Idle Time fields contain valid values, take no actions during configured times.
  - If Idle Time fields do *not* contain valid values, proceed to the next step.
- check and react to configured *Initial Wait Time*
  - If database still unavailable at the end of *Initial Wait Time*, *NOTIFY*
- check and react to configured *Additional Wait Time*
  - If database still unavailable at the end of *Additional Wait Time*, *NOTIFY*
- check and react to configured *Actions*
  - What is done from this point on varies by the Action selected (see *8.1.3.4*, below, for options).

    Options range from *Notify Only* to rebooting the operating system. As illustrated in *Figure 8-1*, *NOTIFY* may occur at several stages during an Action sequence.

8.1.3.1 Purpose of Idle Time

Even when a server is up, and the target database has not been corrupted, successfully opening it can take prohibitively long during times of unusually high application activity. This might be due to scheduled replication, to name but one possibility.

To avoid receiving notifications that a server has (or may have) crashed, Crash Detection provides two options:

- configuring *Idle Time*
  - Crash Detection assumes the database will be unavailable, and takes no actions if it cannot be opened during configured *Idle Time*.
- increasing *Wait Time* parameters
  - Crash Detection allows more time for retries to succeed before deciding a server may have crashed.

Each method has advantages and disadvantages. Take the peculiar characteristics of your environment into account, when deciding which method to employ.

Using *Idle Time* settings to effectively turn off Crash Detection during certain hours has two drawbacks:

- Crash Detection does not react to *actual* server crashes between the times entered in *Begin Idle Time* and *End Idle Time*.
- *Idle Time* settings apply to all days of the week. They cannot be configured to apply on Fridays, and not Tuesdays, for instance.

In short, Crash Detection does not make a determination as to *why* the target database is unavailable. During configured *Idle Time*, Crash Detection ignores the condition and takes no actions.

Increasing the parameters *Initial Wait Time* and *Additional Wait Time* (see *Figure 8.1.3.2* and *Figure 8.1.3.3*, below) has one clear disadvantage.

Crash Detection references these settings at *all times of day*, not just when the server is particularly busy.

Example 1: Initial wait time

Assume that an Initial Wait Time of 120 seconds, along with an Additional Wait Time setting of 180 seconds, would always allow Crash Detection enough time to open the target database.

Consider whether you want Crash Detection *always* to wait five minutes (300 seconds) before concluding that the local server has crashed (and taking actions to notify you/correct the condition).

Before deciding which method is best for a particular server, consider how much longer *Wait Times* would have to be, to handle periods of unusually high database application activity.

Remember ...

Crash Detection works locally, and configuration options can vary from one server to the next, if this best serves the needs of your environment.

TO CONFIGURE IDLE TIME:

1 Click on the Parameter Configuration toolbar icon.

This brings up the interface displayed below.

2 Use the combo box at the upper left to select the server.

3 With the correct server specified, click on the icon (to the right of the combo box).

This brings up a list of folders *for installed components*.

4 Open the Monitor folder.

5 Click on Crash Detection, to access

6 In the right-hand pane, click on the *Begin Idle time* parameter.

The background of the parameter changes from white to blue, and the edit icon is activated.

The parameter is not yet editable. To launch the Edit Parameter dialog, proceed with the following Step.

7 To change the parameter's value, click on the icon; alternatively, double-click the parameter in the right-hand pane of the Console.

An Edit Parameter dialog is displayed, allowing you to enter a new value/values. For details, see *11.1.0.0*.

8 Enter the new value and click OK.

9 Repeat the process for *End Idle Time*.

*In the Server field of the dialog, leaving 'Local' (the default) will change settings only for the server whose **iwparam.nsf** was accessed.*

Enter a comma-delimited list of servers/groups to change the parameter on other systems.

All Crash Detection parameters discussed below are accessed using the same method discussed above.

8.1.3.2 Initial Wait Time

Even when you expect a server to be available, several factors can cause a delay in the *DbOpen* process. *Initial Wait Time* allows you to configure the length of time Crash Detection will wait before it assumes a server *may* have crashed.

As illustrated in *Figure 8-1*, Crash Detection takes no actions during the Initial Wait Time; in fact, it continues to try to open the database.

If Crash Detection is able to open the database before the Initial Wait Time has expired, no steps will be taken, and the program returns and awaits the next monitoring cycle.

If Crash Detection is *unable* to open the database before the end of the configured Initial Wait Time, configured (notification/launch program) actions are taken.

At this point, Crash Detection notifications indicate that the server may have crashed (not that it actually has crashed).

Assuming the database still cannot be opened, Crash Detection will proceed to the next level, namely *Additional Wait Time*.

8.1.3.3 Additional Wait Time

As with Initial Wait Time, no actions are taken for its duration, with the exception of periodic checks made to see if the target database has become available.

Once this second waiting period expires, Crash Detection assumes that the server actually *has* crashed, and configured notifications contain a message to that effect.

8.1.3.4 Recycle/Reboot/Notify?

Once Crash Detection has determined that a server has crashed, the next level of actions to be taken offers four options:

- Notify only
  - Configured notification/launch program actions are taken, but no attempt is made to recycle the server (at either the Notes or the OS level).
- Recycle, do not reboot
  - Configured notification/launch program actions are taken, followed by an attempt to recycle the Domino server.

    However, if the server cannot be recycled, the operating system is not rebooted.
- Recycle, reboot if necessary
  - Configured notification/launch program actions are taken, followed by an attempt to recycle the server at the Notes level.

    If the Domino server cannot be recycled, the operating system is rebooted.

- Reboot
  - Configured notification/launch program actions are taken, followed by a reboot of the operating system.

8.1.4.0 Notification/launch program options

Crash Detection sends takes configured actions when it determines that the Domino server may have (or has) crashed.

For that reason, notification options do not include e-mail, since the local server is deemed unavailable and therefore unable to route mail messages.

Available notification options include:

- Paging
- SNMP traps
- non-Notes notifications launched by proprietary programs Crash Detection is configured to launch

(See also *8.1.5.1*, below.)

8.1.5.0 Other configuration avenues

In addition to the Parameter Configuration Utility, Crash Detection settings can be accessed/modified in the following places:

- Crash Detection UI (NT only)
- registry (NT)

```
HKEY_LOCAL_MACHINE/SOFTWARE/Candle/Inte
  lliWatch/Monitor/Crash Detection
```

- iwmon.ini (UNIX)

```
[Crash Detection]

[relevant setting]
```

- iwparam.nsf

  Before they will take effect, changes made in this database must subsequently be replicated to the servers in question.

8.1.5.1 Using the Crash Detection UI

On NT servers, Crash Detection can be configured on the local machine using the UI (dark-blue icon, located in the right-hand pan of the status bar).

Settings are organized on the following tabs:

- Monitoring
  - database monitored (usually log.nsf)
  - use full path to server

    Selecting the checkbox will cause the DbOpen operation to be done via the server, rather than locally.
- Server Load Adjustment
  - Initial Wait Time
  - Additional Wait Time
- Recycle Options
  - Terminate Notes Time

    Applicable when *Recycle, reboot if necessary* is the selected Action, for instance.
  - Wait for Server to Restart

    Applicable when *Recycle, reboot if necessary* is the selected Action, for instance.
  - Notification level

    Options are:
    - NONE
    - LIMITED: notifies of server successfully restarted
    - VERBOSE: notifies at all significant stages of recycle process

Do not use the Verbose setting unless instructed to do so by IntelliWatch Customer Support, since it produces a significant number of notifications.

  - send page (yes, if checked)
  - send SNMP trap (yes, if checked)
- Actions
  - Notify only
  - Recycle, do not reboot
  - Recycle, reboot if necessary
  - Reboot
- Notification
  - suspects Notes has crashed
    - send SNMP trap (yes, if checked)
    - send page (yes, if checked)
    - run program
  - knows Notes has crashed
    - send SNMP trap (yes, if checked)
    - send page (yes, if checked)
    - run program
- Notes Programs
  - [to be implemented]

8.1.6.0 Wait Time adjustment (NT)

Crash Detection (NT only) dynamically adjusts wait times, enabling the program to adapt to your environment.

8.1.6.1 What adjusts, and what doesn't

As explained under *8.1.3.2* and *8.1.3.3*, there is both an *Initial Wait Time* and an *Additional Wait Time*. Of these, only the *former* adjusts to connectivity history.

Example 2: Wait time adjustments

ServerABC's *Wait Time* settings are:

- *Initial Wait Time = 40*
- *Additional Wait Time = 60*

Assume further that Crash Detection is only able to open **database.nsf** on ServerABC after the expiration of the *Initial Wait Time*, and thirty seconds into the *Additional Wait Time*.

The *Initial Wait Time* setting in the registry will be adjusted upward by thirty (number of seconds of *Additional Wait Time* required to open the database). After this adjustment, ServerABC's new settings will be:

- *Initial Wait Time = 70 [40 plus 30]*
- *Additional Wait Time = 60*

If a future Crash Detection checking cycle takes the full 70 seconds of *Initial Wait Time*, plus twenty seconds of *Additional Wait Time*, the settings will be further adjusted, as follows:

- *Initial Wait Time = 90 [70 plus 20]*
- *Additional Wait Time = 60*

The Additional Wait Time setting does not change.

The Initial Wait Time parameter value does not reset when the server is restarted, since any adjustments are made in NT registry settings.

Initial Wait Time can be reset manually, using the Parameter Configuration Utility, or one of the other access methods detailed under 8.1.5.0 and 8.1.5.1.

This example refers only to wait-time adjustments on NT, since the program functions somewhat differently on UNIX systems.

FIGURE 8-1: Crash Detection program flow

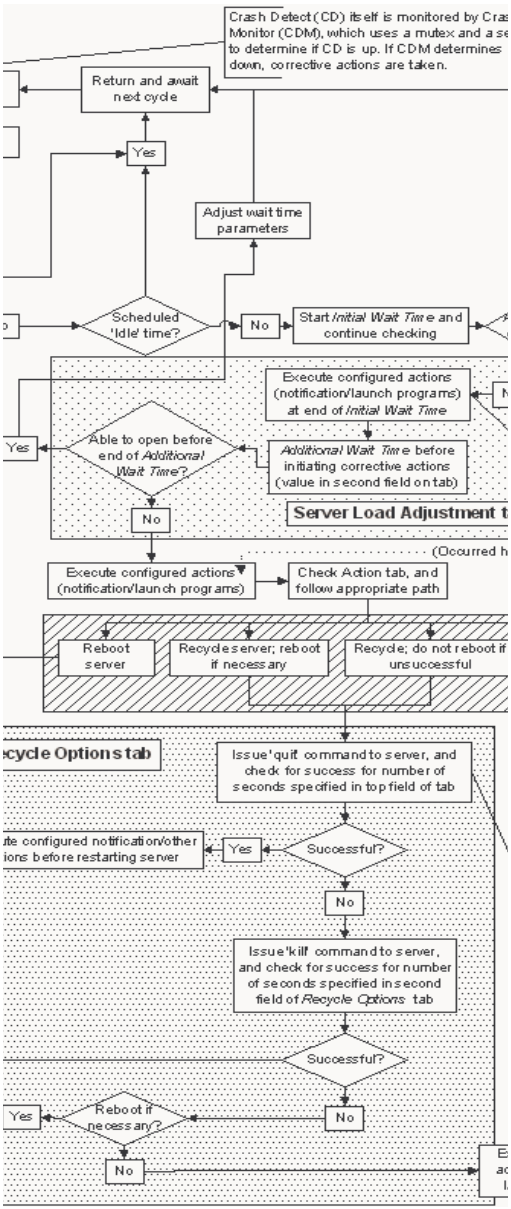

Chapter

Analyzer

9

IntelliWatch Analyzer is the report engine for IntelliWatch. Analyzer retrieves both IntelliWatch and native Notes statistics, then turns the data into HTML reports for e-mail distribution.

In addition, statistics generated by Advanced ServerWatch Activity Logging can be retrieved for incorporation in reports (see *Figure 9-1 on page 290*).

Chapter Contents

FIGURE 9-1: IntelliWatch Statistic Architecture

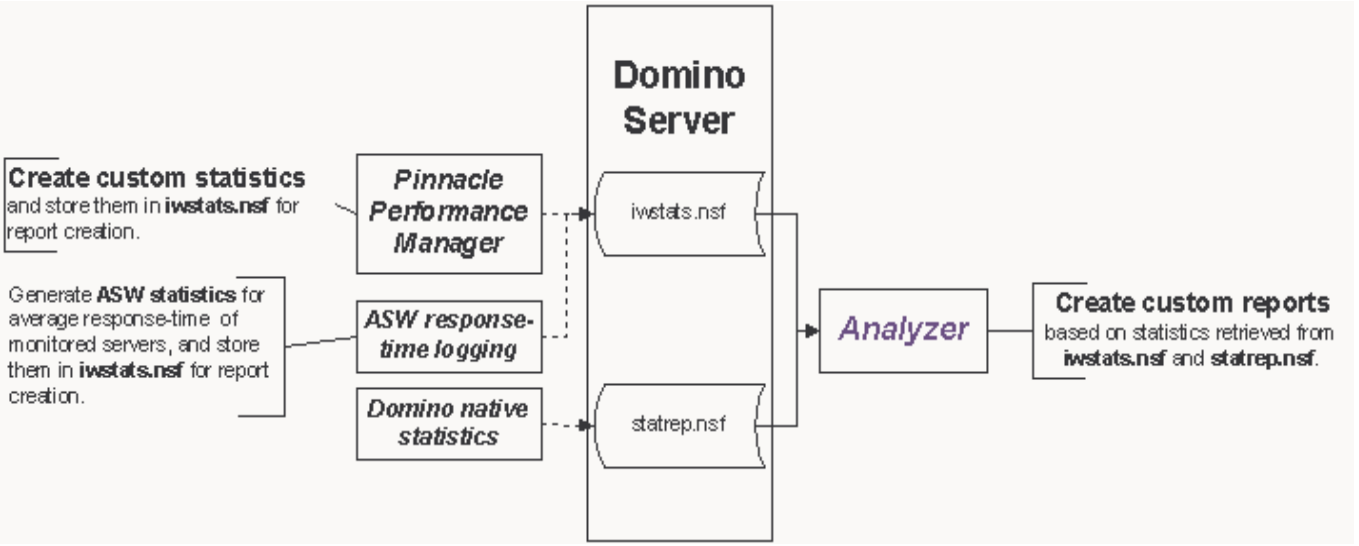

9.1.0.0 OVERVIEW

As illustrated in *Figure 9-1*, Analyzer is responsible for ***retrieval*** and ***customizing*** of data generated by PM

Before this can occur, however, the following steps must have been taken:

- Creation of PM statistic definitions in **iwpmstat.nsf** (and replication of those definitions throughout your environment).
- Generation of data for enabled statistics (by the IWSTATG task, assuming the relevant events, etc., have occurred).
- Replication of this data to statistic databases on the Analyzer server.

Before discussing details of statistic generation and configuration, a few words are in order about the Analyzer server task.

9.1.1.0 Analyzer Server

Analyzer runs as a Domino server task (named **analyzer**).

As a rule, you'll have only one or two instances of Analyzer Server running in your environment.

9.1.1.1 When does the task run?

Analyzer Server installs on your statistics-collection server(s) *only*, and runs:

- on a schedule
  - The launch time for the Analyzer task is specified in the **notes.ini**, and defaults to 3AM (**ServerTasksAt3=Analyzer**).

    Modify the time (if so desired) by changing the digit just preceding the equals sign. Use the 24-hour time format. (To launch Analyzer at 8pm, for example, modify the line to read **ServerTasksAt20=Analyzer**.
- when launched manually

  Type:

```
load analyzer [name of report]
```

  at the Server Console, or

```
analyzer [name of report]
```

  at a command prompt.

To run a single scheduled report, include its name on the command line; to run all *scheduled reports, leave off the last argument ([name of report]).*

Generating a selective list of reports (comma-delimited, for example) is not supported.

For more details on report scheduling, see *"Scheduling" on page 304*.

9.1.2.0 Analyzer data sources

Analyzer reports on three types of statistics:

- IntelliWatch statistics
- ASW activity logging
- native Notes statistics

The first two statistic types are retrieved from **iwstats.nsf**; native Notes statistics are retrieved from **statrep.nsf**.

The default location for these databases on R5 (on NT) is:

- statrep.nsf:

  /Lotus/Domino/Data
- iwstats.nsf:

  /Lotus/Domino/Data/IntelliWatch

9.1.3.0 Analyzer Configuration

All configuration is done via the **analyzer.nsf** database, located on your statistics-collection server. Located in the same folder as all other IntelliWatch databases, the default location on R5 is: /Lotus/Domino/Data/IntelliWatch).

The Analyzer configuration database can be accessed three ways:

- **Pinnacle Console via a browser**
- **via the stand-alone IntelliWatch client**
- **via a Notes client**

9.1.4.0 Location of the Analyzer Server

The Primary Server must be configured with the location of the system running Analyzer Server (seldom are they the same).

Although the Setup gives you the option of not installing Analyzer Server on a given system, it does *not* allow you to tell that system where Analyzer will be running.

Since errors can result, after installing IntelliWatch on any system to which you will connect with the Console—but that will not be running Analyzer Server—take the following steps.

The following procedure assumes you've already succeeded in accessing the Pinnacle Console (either via the web, or using the stand-alone version).

TO ENTER A LOCATION FOR ANALYZER SERVER:

1. Click on the  toolbar icon; alternatively, go to **Solutions > Parameter Configuration** via the drop-down menus.
2. Make sure the name of the target system appears in the combo box on the toolbar.
3. Click on the  toolbar icon.
4. Once configuration settings are loaded in the left-hand pane, open the **Analyzer** folder.
5. Again in the left-hand pane, select the value **Current State**.
6. In the right-hand pane, select the parameter **ServerList** and click on the  toolbar icon (or double-click the parameter).
7. In the pop-up dialog, enter the name of the system running Analyzer Server (can be a comma-delimited list, for multiple instances of Analyzer).
8. Click OK.
9. Exit the Console, then relaunch it.
10. Now, when you click on the  toolbar icon, the folder in the left-hand pane will bear the name of the Analyzer Server.

The above procedure changes the relevant setting ONLY on the Primary Server.

9.1.5.0 Creating Contacts

When accessing the Analyzer database over the Web, Servers, Groups, and Persons that you want to appear in selection lists must first have a Contact document created for them in **iwparam.nsf** on the Analyzer Server. (The default is *, or all servers.)

A feature of Contacts is that when statistics, charts and reports are being created, only those Contacts relevant to the configuration procedure being carried out are displayed in the list box. For example, no Persons are displayed in a Server Name field, and no Servers are displayed in a Send To field.

Contacts must be created via a Notes client, and come into play only when Analyzer is being configured over the Web.

TO CREATE CONTACTS IN IWPARAM.NSF:

1. Open **iwparam.nsf** by going to **File > Database > Open > [Name of Analyzer server] > IntelliWatch folder > etc.**
2. Go to **Create > Contact** via the drop-down menus.
3. Select Server, Group or Person via the radio buttons at the upper left.

4 Click on the down arrow of the combo box to launch a selection dialog for the chosen Contact type.

5 Select the desired Contact, then click OK.

If the desired Contact is not found, it can be filled in by hand in the field at the bottom of the dialog. Click OK to complete the addition.

6 Press Esc, which pops up a confirmation dialog.

7 Click Yes (the default) to save the document.

9.2.0.0 REPORTING GUIDELINES

Please review *Figure 9-2*, as well as the Reporting Guidelines that follow, before creating reports.

FIGURE 9-2: Preparing to create a report in Analyzer

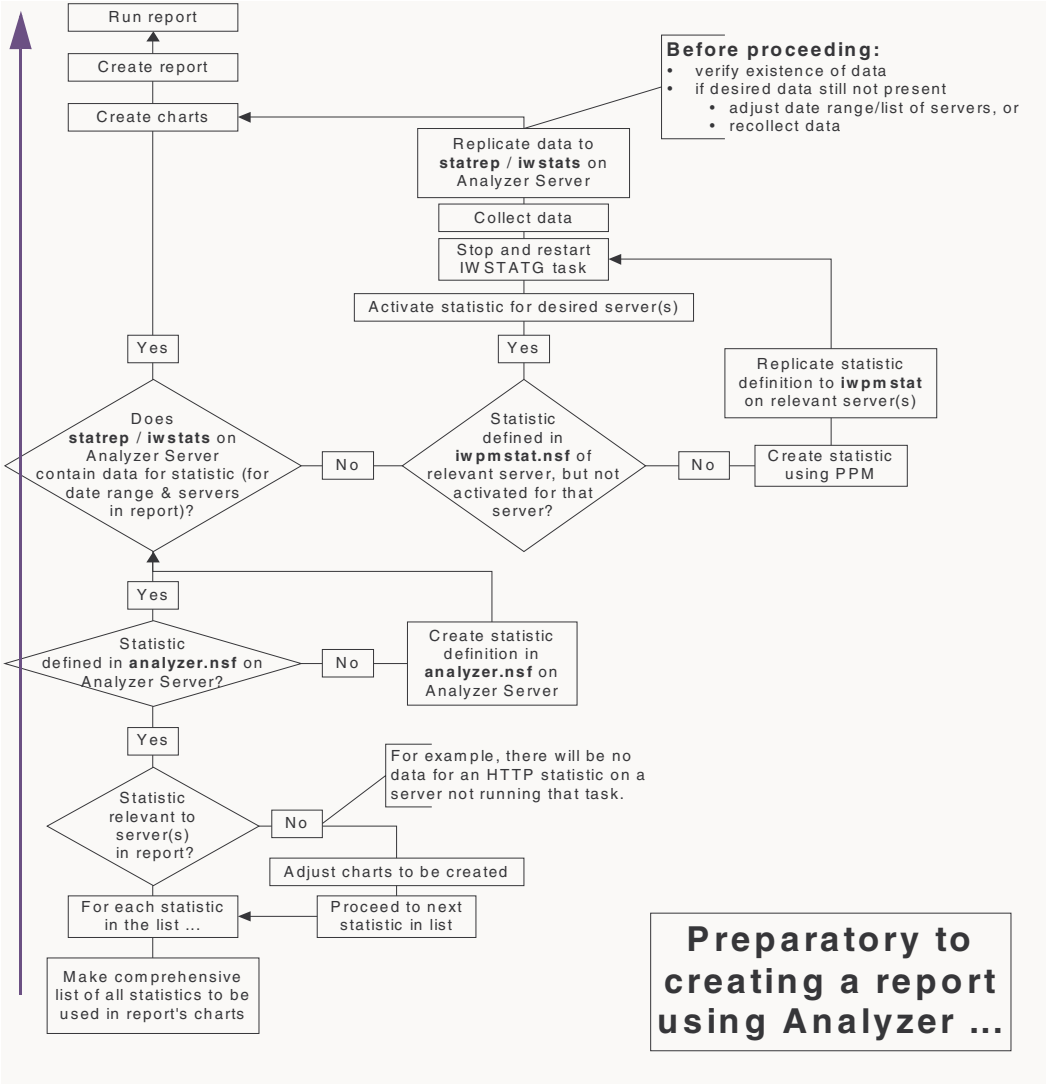

Bear in mind the following guidelines, to save time, and to produce more effective reports.

- **Analyzer vs PM statistics**
  - Understand the difference between PM and Analyzer statistics
- **Common reasons for report failure**
  - Procedural order
  - Configuration errors
- **Start at the bottom**
  - Start at the statistic level, when creating Analyzer reports
- **Collect only meaningful Data**
  - Be selective when collecting data, to save disk space and processing time
- **Names are important**
  - Name statistics in ways that make their purpose clear at a glance
- **Wildcard usage**
  - Server Lists
  - Statistic Names

9.2.1.0 Analyzer vs PM statistics

The difference between statistic creation in PM, and statistic definition in Analyzer, can be summed up in two simple sentences:

- PM statistics *generate* data (by type).
- Analyzer statistics *retrieve* data (based on the selection criteria you build into them).

PM statistic definitions are not inherited by Analyzer for two reasons:

- PM statistics are generative, whereas Analyzer definitions are for retrieving data.
- This design allow you to create Analyzer statistic definitions that permit selective retrieval of PM data.

9.2.1.1 Generating the data

The following example is based on the PM statistic type *Mail Domain*, which type returns the following four data elements (for each domain for which there is mail traffic):

[Statistic Name].incoming.[domain name].Count
[Statistic Name].incoming.[domain name].Size
[Statistic Name].outgoing.[domain name].Count
[Statistic Name].outgoing.[domain name].Size

For the following examples, we'll call our sample statistic Mail.Domain.

Assuming there was mail traffic between your organization and the two Internet domains *hotmail.com* and *npr.org*, the following data elements will be written to **iwstats.nsf** (for statistics enabled for Reporting):

- Count (number of messages by domain, both incoming and outgoing)

Mail.Domain.hotmail.com.incoming.Count=[Value]
Mail.Domain.hotmail.com.outgoing.Count =[Value]
Mail.Domain.npr.org.incoming.Count =[Value]
Mail.Domain.npr.org.outgoing.Count =[Value]

- Size (total bytes by domain and direction):

Mail.Domain.hotmail.com.incoming.Size =[Value]
Mail.Domain.hotmail.com.outgoing.Size =[Value]
Mail.Domain.npr.org.incoming.Size =[Value]

Mail.Domain.npr.org.outgoing.Size =[Value]

9.2.1.2 Retrieving the data

Analyzer allows you to retrieve any—or all—of the above data elements, simply by creating the correct Analyzer statistic definition. Continuing with the example under *9.2.1.1*, here are a few examples of how to retrieve various components of the above data:

- all stats for hotmail.com

Mail.Domain.h*

- Counts ONLY for npr.org:

Mail.Domain.n*.Count

- Size elements for both domains:

Mail.Domain.*.Size

- All data for both domains

Mail.Domain.*

9.2.2.0 Reasons for report failure

There are several reasons why Analyzer reports may fail, but virtually all of them relate to errors, either of the order in which procedures are performed, or mistakes in general configuration:

9.2.2.1 Errors in procedural order

Analyzer reports will fail if you run them when:

- the Analyzer statistic definitions in the report have no corresponding PM statistic
- an Analyzer statistic definition has a corresponding PM statistic, but no data has yet been generated for the latter
- the data you want to retrieve has not yet been replicated to the server running the Analyzer server task

9.2.2.2 Errors in statistic configuration

The two most common configuration errors are related to:

- a lack of data for the time period, statistic, or server to be included in the report
- failure to consider what the PM statistic returns

Take the statistic we created under *9.2.1.1*. Based on the assumptions we made, you'll recall that the statistic generated a specific set of data for the two Internet domains for which there was mail traffic.

But here's the rub!

Although the PM statistic itself was simply called *Mail.Domain*, the IWSTATG server task adds other elements to that base name at the time of data generation (what it adds depends on the PM type—for a few data types, nothing is added).

To emphasize what PM added to the Statistic Name in the present case, it's now in orange:

- Count

Mail.Domain.hotmail.com.incoming.Count=[Value]
Mail.Domain.hotmail.com.outgoing.Count =[Value]
Mail.Domain.npr.org.incoming.Count =[Value]
Mail.Domain.npr.org.outgoing.Count =[Value]

- Size

Mail.Domain.hotmail.com.incoming.Size =[Value]
Mail.Domain.hotmail.com.outgoing.Size =[Value]
Mail.Domain.npr.org.incoming.Size =[Value]
Mail.Domain.npr.org.outgoing.Size =[Value]

Although all you needed to create in PM was a statistic of Type Mail Domain*—to generate ALL the above data—trying to retrieve the above data using an Analyzer statistic definition called simply Mail.Domain will fail.*

We suggest you review the sample Analyzer definitions under *9.2.1.2*.

Which brings us to perhaps the most fundamental guideline of report creation.

9.2.3.0 Start at the bottom

'Starting at the bottom' can be broken down into the following three rules:

- *Having* statistical data is more important than *defining* it.
- Reports cannot incorporate charts that don't yet exist.
- Charts cannot incorporate statistics that are not yet defined.

But what of the apparent contradiction between rules one and three?

Rule one doesn't say defining statistics isn't important, only that having the corresponding data is even more important.

What does that mean in practice?

Example 1: Generating a report

As long as **statrep.nsf** and **iwstats.nsf** (on the Analyzer server) already contain the statistical data you want, report creation is as simple as following these steps.

TO GENERATE A REPORT:

1. Define Analyzer statistics corresponding to the stored data.
2. Create charts incorporating those statistics.
3. Include those charts in a report.
4. Run the report.

If, on the other hand, the data you want *aren't* contained in **statrep.nsf** or **iwstats.nsf** on the Analyzer server (and can't be replicated to that server from somewhere else), Analyzer statistic definitions are irrelevant, and chart production fails.

9.2.4.0 Collect only meaningful data

Not all data are meaningful, so establish that a given statistic is meaningful in the context of your environment before you collect it.

By pre-selecting the data to be collected, you can potentially:

- save valuable disk space
- enhance the readability of the list of collected statistics by limiting their number
- improve Analyzer Server performance by limiting the amount of data being processed

Example 2: Relevance to server

The value of the statistic DiskC.Free may be quite relevant on certain servers, but not on others. Or, at a minimum, you made need to check this statistic only daily on some servers, while on others it should be checked as often as hourly (or as seldom as weekly).

Check whether servers are organized logically into groups in your NAB. If so, create a version of DiskC.Free appropriate to each group. If not, consult your Notes documentation for advice on how best to group Domino servers.

Finally, let us consider the value of good statistic naming practices.

9.2.5.0 Names are important

PM allows you to name statistics in ways reflective of their role in your environment—whether based on native Notes statistics, or NT Performance Counters.

Here's a practical example of how statistical names can be made to represent the data being collected.

Example 3: Remote server needs

You want to keep closer track of free disk space on your remote servers than on servers at the home office, doing so on an hourly basis.

While Notes includes the statistic DiskC.Free, this name is not indicative of the location of these servers in your company.

In PM (and subsequently in Analyzer), define a statistic called

- RemoteSvrs.DiskC.Free.Hourly

For purposes of data collection, associate the statistic with the group RemoteSvrs.

Once these data are collected and replicated to your Analyzer server, creating charts incorporating this statistic is made easier by the fact that the statistic name is clearly representative of the data to be contained in the report.

9.2.6.0 Wildcard Usage

Perhaps the most important 'secret' to creating Analyzer reports is wildcard usage.

Wildcards are a two-edged sword, however. On the one hand, they allow you almost infinite customizing in retrieving data. At the same time, they can cause much more data to be retrieved than is required for a given report.

When wildcards are also used to generate the data (in PM)—and this is particularly true of PM's Mail types—such a large number of unique statistics can be generated that performance (and even functionality) are affected.

9.2.6.1 Basics

Three wildcards are supported:

- # — for any numeric character
- ? — for any single character
- * — for zero or more characters

9.2.6.2 Server Names

Assume for the moment that the names of all servers in your environment are identical (MySvr[numerical character]/MyDomain),

with the exception of the single numeric character here shown in brackets.

MySvr#/MyDomain

would retrieve available data for all servers.

If, on the other hand, you wanted to retrieve data for all servers in your company's /Northern/ domain, simply use the string "*/Northern/*".

Remember ...

to consider the factors that can affect data display (see "One to Many" on page 300 and "Data display and legend size" on page 301).

9.2.6.3 Statistic Names

Wildcard usage in Statistic Names is even more critical than in Server Names (after all, your servers will most likely be in groups, and you may very well want to retrieve a given statistic for all members of a group).

With statistics, however—and especially those statistics where PM adds one or more elements to the PM statistic name—wildcards may be the only practicable means of retrieving all but the smallest quantities of data.

Consider the example of the PM Replication statistic type. Assuming you gave it the name *Rep.Stat*, one of the elements PM returns is:

Rep.Stat.Pulled.KBReceived.[server name]

Without wildcards, you would have to create an Analyzer statistic definition that included the entire server name!

By simply substituting a "*" for the server name, you could retrieve data for all servers that had relevant replication activity (vis à vis the server on which the data was generated).

By the same token, if too many servers fit the criteria, you may have to use a mixture of wildcards and characters, to avoid display problems (see "Labels and the number of statistics reported:" on page 312).

9.3.0.0 ANALYZER DIALOGS (COMMON FIELDS)

The section *"Start at the bottom" on page 297* explains the importance of beginning report building at the statistic level. Dialog fields are therefore discussed from the lowest to the highest level, starting with Statistics and ending with Reports *on pages 308 to 322*.

9.4.0.0 WORKING WITH STATISTICS

All procedures outlined below assume you have already accessed the Analyzer interface at the Pinnacle Console.

9.4.1.0 Defining Statistics

Before IntelliWatch statistics can be retrieved by Analyzer server, they must be associated with statistic definitions in **analyzer.nsf**.

Statistic definition is a straightforward process that is the same for all statistic types.

TO DEFINE ANALYZER STATISTICS:

1 Go to **Options > New > Statistic** via the drop-down menus; alternatively, click on the toolbar icon.

2 Fill in the Title of the statistic.

Use front-loaded Titles (Names) to facilitate statistic selection.

3 Fill in the Statistic field using the Statistic Selection combo box; alternatively, type in the desired value.

Many—if not most—Analyzer statistic definitions will include wildcards, so select the PM statistic first, then edit the value in the field.

4 Using the radio buttons, select the database where the corresponding PM/Notes statistic resides.

5 Fill in the For Server field using the Server Selection dialog; alternatively, type in the name of the server(s) to which the statistic applies.

For wildcard usage for both Server and Statistic Names, see "Wildcard Usage" on page 298.

6 Save the statistic definition by going to **Options > Save** via the drop-down menus; alternatively, click on the toolbar icon.

9.4.2.0 Creating Charts

Once Analyzer statistic definitions have been associated with available IntelliWatch stats, they can be incorporated into charts.

Unlike statistic definition, which is the same for all statistic types, the steps for chart creation can vary widely, depending on the options selected. For details on chart options, refer to the usage information *on pages 310 to 317*.

9.4.2.1 One to Many

Before you begin creating charts, bear in mind that Servers and Statistics are in a one-to-many relationship. What does this mean? Simply that a chart can contain *either*:

- **multiple statistics for one server** (a maximum of 16 statistics is strongly recommended, to avoid potential display issues)
- **multiple servers for one statistic** (again, no more than 16 servers should be included in a single chart, to avoid potential display issues)

9.4.2.2 Wildcard usage

Wildcard usage in charts is likewise based on the *one-to-many* principle. Wildcards can be used *either*

- in a chart's Server List, *or*
- in a chart's Y-Axis Statistic List(s)

but not both.

Put another way, you can have *either*

- a chart for a single statistic, with wildcards in the Server List, *or*
- a chart for a single server, with wildcards in the Y-Axis Statistic List(s)

As discussed under *9.4.2.1*, above, display problems can occur when too many statistics (or servers) are displayed on a chart. Wildcards in charts can lead to just such display issues.

9.4.3.0 Data display and legend size

Analyzer balances the relative size of the data display and the legend, giving precedence to the latter.

Analyzer attempts to create a legend incorporating all charted items. As the number of individual lines of data to be displayed increases, the legend is *enlarged* to accommodate them; the data display is commensurately *reduced* in size, eventually to the point that chart readability can suffer.

(Giving precedence to the display field rather than to the legend might, at first, seem a more logical approach. However, this would eventually lead to the legend's becoming so small as to be unreadable. A chart with clearly legible data, but an unreadable legend, is no more useful than the reverse.)

9.4.3.1 A rule of thumb...

Whenever wildcard usage would result in more than 10-12 lines of data, adjust wildcarded strings to reduce data lines displayed to that number (or less).

Example 4: Too many servers

Thinking you have only ten servers in your environment beginning with the letter "M", you use "M*" in the Servers field in a chart tracking *Disk.C.Free*. Unbeknownst to you, however, twelve other servers beginning with that letter have recently been added to the environment.

When the chart is generated, there are twenty-two lines of data to be displayed, rather than ten. All twenty-two are listed in the legend, leaving less room for data display. You may find that readability of the displayed data suffers to an unacceptable degree.

Example 5: Too many statistics

You want to chart all statistics associated with CPU utilization on the server *HQSpoke3*. (Your Notes team names such statistics *HQSpoke3.CPU.[Process]*.)

To your knowledge, the names of no other statistics collected on this server begin with the letter "C" (after the server name), so you enter *HQSpoke3.C** in the (Left) Y-Axis Statistic field of the new chart.

Recently, however, your Notes diagnostic team has added several CPU statistics for troubleshooting purposes (calling them *HQSpoke3.CPU_Test.[Process]*.

When Analyzer creates the chart with the wildcard usage you employed, all these "**.Test.**" statistics will be included, and the chart data display will likely be unsatisfactory.

A simple adjustment in the wildcarded string gives you the chart you want to see: *HQSpoke3.CPU.**

Now, none of the *.CPU_Test.** will be included in the chart.

The following steps are generic. For option-related details, see the usage instructions for individual dialogs on the preceding pages.

TO CREATE ANALYZER CHARTS:

BASIC INFORMATION TAB

1 Go to **Options > New > Chart** via the drop-down menus; alternatively, click on the  toolbar icon.

2 Fill in the Title so as to be able to see at a glance what the chart represents.

3 Fill in the Chart Name (or filename for the generated .gif).

Remember ...

This is a filename, not a Chart Title. Use naming conventions appropriate to your environment.

4 Enter a summary of the chart's purpose, to be published beneath the data display.

5 Select the chart type.

6 Decide if you want to select any of the Optional Advanced Settings. If so, check *Enabled*, along with the desired option.

7 Select the analysis (data) to be displayed.

8 Give the x- (horizontal) axis a label reflective of what is to be displayed.

9 Select the time scale of the x-axis.

10 To save your changes before proceeding, go to **Options > Save** via the drop-down menus; alternatively, click on the  toolbar icon.

LEFT-AXIS INFORMATION

1 Enter a name for the left-hand y- (vertical) axis, that represents the data being displayed.

2 Use the Select Statistic dialog to associate statistic(s) with the left y axis.

3 Select the number of decimal points to display (the default is 2).

4 Select the appropriate Range Type for the data to be displayed.

5 Select the desired Range, or units of division.

Not visible for the Percent option.

6 To save your changes before proceeding, go to **Options > Save** via the drop-down menus; alternatively, click on the  toolbar icon.

RIGHT-AXIS INFORMATION

1 Enter a name for the right-hand y- (vertical) axis, that represents the data being displayed.

2 Use the Select Statistic dialog to associate statistic(s) with the right y-axis.

3 Select the number of decimal points to display (the default is 2).

4 Select the appropriate Range Type for the data to be displayed.

5 Select the desired Range, or units of division.

This option is not visible for the Range Type Percent.

6 To save your changes before proceeding, go to **Options > Save** via the drop-down menus; alternatively, click on the  toolbar icon.

ADVANCED INFORMATION

1 Select the Create CSV files option if you want to export retrieved data as comma-delimited data files.

2

3 Use Report Settings is the recommended setting for the Period Covered option, unless the chart's date range settings should differ from those of the report.

4 To save your changes, go to **Options > Save** via the drop-down menus; alternatively, click on the toolbar icon.

9.4.4.0 Creating Reports

After all desired charts have been created, report creation can commence.

TO CREATE ANALYZER REPORTS:

BASIC INFORMATION

1 Go to **Options > New > Report** via the drop-down menus; alternatively, click on the toolbar icon.

2 Fill in the Title so as to be able to see at a glance the Report's purpose.

3 Fill in the Report Name (or filename for the generated .html).

Remember ...

This is a filename, not a Report Title. Use naming conventions appropriate to your environment.

4 Enter a Report summary.

5 Using the combo box, select the time interval at which to run the report.

For more information on the options, see *page 319*.

6 Using the combo box, select the day on which to run the report.

7 Fill in the Mail Report To field using the Select Mail Addresses dialog; alternatively, type in a (comma-delimited) list of mail recipients.

If the Mail Report To field is left blank, the configured IntelliWatch parameters determine report recipients.

8 To save your changes before proceeding, go to **Options > Save** via the drop-down menus; alternatively, click on the toolbar icon.

DATA INFORMATION

1 Use the Select Charts dialog to populate the list box of charts to include with the report.

For usage information on the Select Override Chart Settings option, see *page 320*.

2 Select the desired date and time settings.

For usage information, see *page 320*.

ADVANCED INFORMATION

The Advanced Information tab governs two ancillary functions:

- CSV (comma-delimited) file creation
  - In addition to Yes and No, the drop down allows you to select the Use Chart Settings option.

In most cases, Use Chart Settings produce the results you want.

Report-level settings take precedence over chart-level settings.

- non-default database locations
  - Notes database name
  - IntelliWatch database name

Leave both these fields blank *unless* **statrep.nsf** or **iwstats.nsf**, respectively, have been given different names, or are in non-standard locations.

  - Report location

Leave this field blank *unless* Analyzer reports are to be stored in a non-standard location.

9.4.5.0 Editing Reports

The procedures for editing reports are the same as those for report creation. Refer to the relevant section to edit a Statistic, Chart or Report.

Remember ...

Always assign a new and unique name to edited Statistics, Charts and Reports.

9.5.0.0 SCHEDULING

Scheduling occurs at two levels:

- Date/Time when report is to run
- Date/Time of retrieved data

Date/Time when report is to run

Report run-times have the following options:

- Daily

This report runs once a day, at the time specified in the Notes.ini (see *"When does the task run?" on page 291*).

- Weekly

This report runs once a week, on the day selected in the combo box, and at the time specified in the Notes.ini.

- nth Day of month

Useful when a report should be run on more than one day of the month, but not on a daily or weekly schedule. Use integers from 1-31 to specify days, with commas as the delimiter. Here too, the report runs at the time specified in the Notes.ini.

- One Time Only

This scheduling option causes the report to be run once, the next time Analyzer Server launches (either on schedule or when manually launched). Scheduling for that report is then changed to Do Not Run.

  - Do Not Run

This scheduling option causes the report to be 'put on hold'. Use this option for reports that are not currently relevant, but may again become so.

Do Not Run also allows you to set up reports in advance: for instance, for servers that are currently being installed, but are not yet in service.

- Period covered
  - Current

    day
    week
    month

Current assumes a normal calendar week, beginning on Sunday and ending on Saturday.

  - Previous

    day
    week
    month

Previous also assumes a normal calendar week, beginning on Sunday and ending on Saturday.

When creating reports, consider carefully the period for which you want to collect data.

A report run on Saturday, 5/17 for the Previous Week, returns data for 5/4 through 5/10.

– Custom

dates

times

Custom must be used whenever:

– Your Work Week is non-standard.

For instance, if your office is only open Tuesday through Friday—*and those are the only days for which you want to retrieve data*.

– You want to incorporate data for only certain days of the week.

Useful if you are suffering occasional resource shortages, but suspect they are limited to certain days of the week. Use Custom as the Period Covered, and troubleshoot by selecting various combinations of days

– You want to incorporate data for only certain times of day.

Useful for troubleshooting issues during peak usage times, for example.

- Dates
  - Applies only when Period Covered is Custom, and includes:

    start date

    end date
- Times
  - Applies only when Period Covered is Custom, and includes:

    start time

    end time
- Work Week
  - Applies only when Period Covered is Custom, and allows the selection of those days that constitute the Work Week for the company (or other organizational unit) in question.

9.6.0.0 REPORT ARCHIVING

Analyzer reports normally run on a schedule, for the previous (or current) day, week, month, and so on.

What does report scheduling have to do with archiving? Simply this: *by design*, whenever Analyzer runs a new 'issue' of a report, it overwrites the previous 'issue'.

Example 6: StatAB_PrevWeek

You have created a report *StatAB_PrevWeek* that runs on the first day of the week, and incorporates Statistics A and B for the previous week.

On Day 1 of Week *2*, data for Week *1* (the previous week) is retrieved and incorporated into the report. On Day 1 of Week *3*, the report imports data for Week *2*—and overwrites *StatAB_PrevWeek*.

Why is the report overwritten? Because it incorporates statistics for *a relative* date range, rather than a *specific* date range. (After all, the report was not entitled *StatAB_Week1[Year]*, but *StatAB_PrevWeek*.) In short, overwriting keeps the report current.

9.6.0.1 Saving individual reports

The report in the above example can be saved such that it will not be overwritten, but

the procedure depends on the form in which you want to archive the information.

- Notes document with attachments
  - Move it to a database of your choosing, and give it a name indicative of the data it contains.

    To save the first of the reports in the previous example, change its name to *StatAB_Week1[Year]*.
- HTML documents
  - Rename the HTML document to *StatAB_Week1[Year]*, and store it—along with the associated gif files—in a location of your choosing.
- CSV file(s)
  - Rename the CSV file *StatAB_Week1[Year]*, and store it in a location of your choosing.

9.6.0.2 *One Time Only* reports

The *one time only* option is especially useful when incorporating statistics for a non-standard date/time range.

Example 7: AB_Last9Mo_AfterHours

As in the previous example, you are interested in Statistics A and B. In this case, however, you want to report on the last nine months, and only on the hours during which the office was closed.

You might call this report *AB_Last9Mo_AfterHours*, for instance.

For details on Scheduling options, see *"Scheduling" on page 304*.

9.7.0.0 Troubleshooting report creation

Occasionally, report creation will fail for reasons not immediately apparent.

Figure 9-3, and the checklist below, were designed to help you determine where and why data retrieval and/or report creation failed.

9.7.1.0 Checklist

- Mis-spelled statistic names
  - Analyzer cannot retrieve data if its statistic definition does not *exactly* agree with what the PM stat returns (including supported wildcard usage).
- incorrect date range
  - **iwstats.nsf** may contain data for a given statistic definition, but if there is no corresponding data for the specified date range, the process fails.
- Data in the wrong database
  - If data are not transferred to the Analyzer Server, the process fails.

9.7.2.0 No .gifs in **iwreport.nsf**

If you want the chart's .gif files attached to reports stored in the database

(**iwreport.nsf**)--as was done in Pinnacle 99--you need to add the following registry key, and set it to 1:

[HKEY_LOCAL_MACHINE\SOFTWARE\Candle\IntelliWatch\Analyzer\Preferences]

"Add Attachments"=dword:00000001

To disable this feature, set the value to 0. This feature is non-standard, and therefore not available via the Setup.

FIGURE 9-3: Analyzer Program Flow

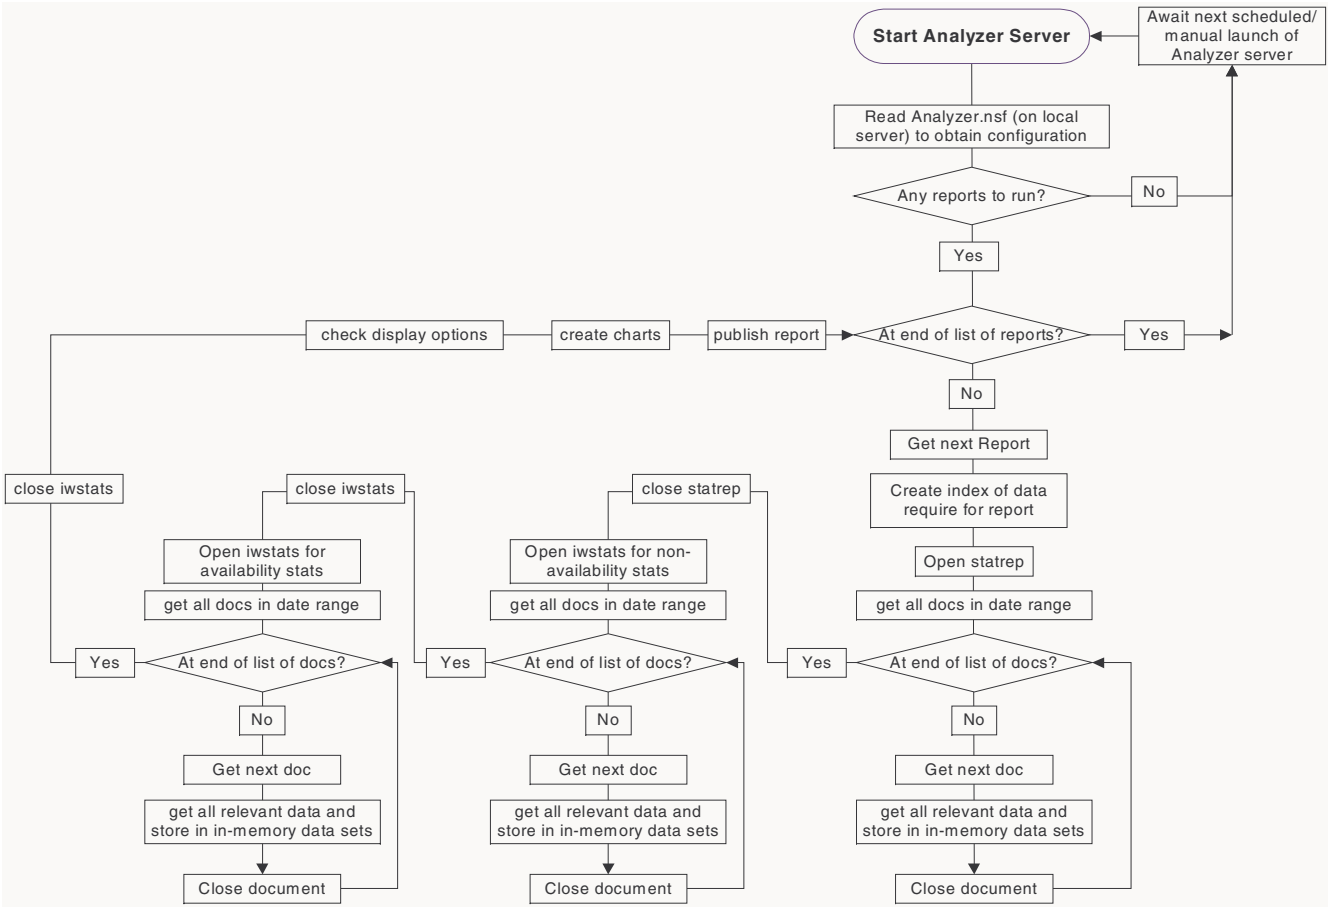

Analyzer Dialogs: Statistic Definition

Basics ...

Statistic definitions:

- Title
  - Specific to Analyzer. A unique identifier used to display available statistic definitions.

    The Title functions independently of the statistic Name, and need not correspond to it in any way.
- Name
  - Must corresponds to all (or part) of what the corresponding PM statistic returns.
- Data Location
  - Specify database in which data resides, either iwstats.nsf for IntelliWatch statistics, or statrep.nsf for native Domino statistics.
- Server(s)
  - Specify which server(s) to include in the data search.

Putting it into practice ...

Statistic Names in Analyzer vs PM:

As discusses under *"Analyzer vs PM statistics" on page 295*, statistic names in Analyzer must correspond to what is returned by their counterparts in PM.

What is returned:

The single most critical consideration when defining statistics in Analyzer is not the *Name* of the PM statistic you want to retrieve, but what that statistic type returns (this includes the PM statistic name, but also—in most cases—additional elements added by the IWSTATG server task).

For details of what is returned by each PM statistic type, the reader is referred to *"Data Returned by PM Statistic Types" on page 459*. For a practical example, based on the *Mail Domain* statistic type, see *"Retrieving the data" on page 296* and *"Errors in statistic configuration" on page 296.*

Analyzer: Statistic Definition

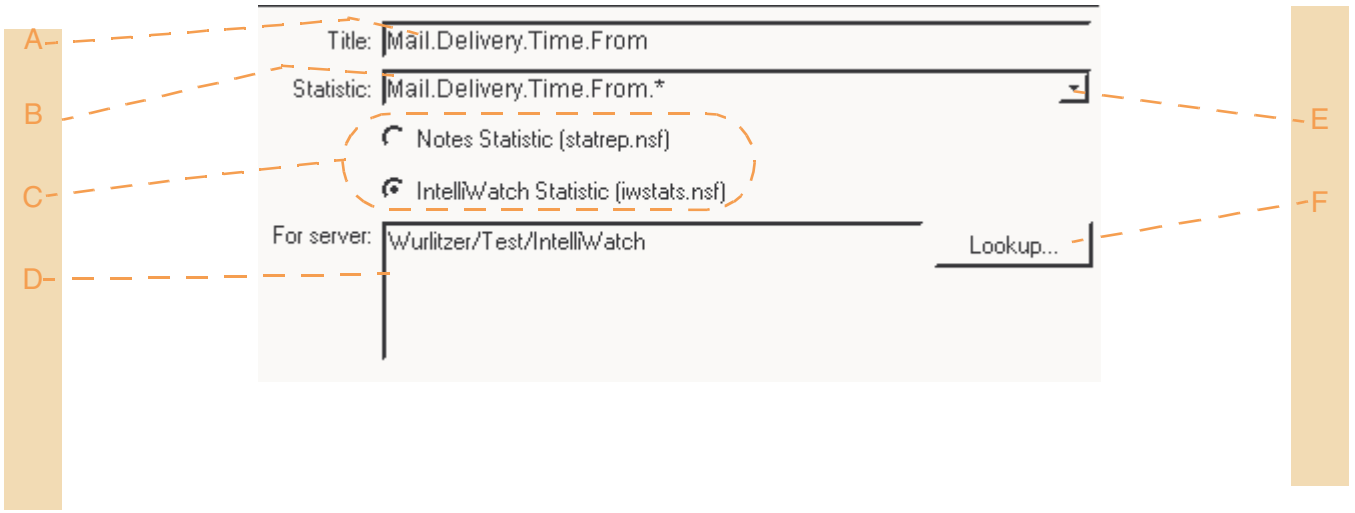

A: title of statistic
B: name of statistic (for applicable naming conventions, see usage instructions, above)
C: statistic repository
D: servers to which server applies
E: combo-box for statistic selection
F: button launches Select Server dialog

Analyzer Dialogs: Chart (Basic Information)

Basics ...

Use this dialog to enter:

- title of chart
- name of .gif file created
  - Illegal characters in chart names: \ / : * ? " < > |
- summary information re the statistics displayed
- optional configuration details
- type of data to present (see Raw vs Average, below)
- x-axis label
- time scale applicable to x-axis

Raw vs Average:

Raw allows incorporation of up to 16 statistics (total for *both* Y-axes), but *only* if all stats are for a single server. Otherwise a maximum of two (2) statistics is allowed (whether you are using one or both Y-axes).

Average allows a maximum of one statistic (*not one per Y-axis*), irrespective of the number of servers associated with it.

Wildcards are allowed in both server lists and statistic lists. For usage details, see *9.4.2.2 on page 300*.

Putting it into practice ...

Title vs Name fields:

The Title appears in two places:

- in the tree view, listing created charts
- at the top of the chart when created as a .gif file

The Chart Name (plus the .gif extension) is the file created and displayed by the report.

Optional Advanced Settings:

To make these options available, check *Enabled*.

- Trigger Threshold
  - Create chart *only when threshold exceeded*.

Example: You want a report on Mail.Dead only when it exceeds Value X. With this option selected, the chart is generated only when there are data points in excess of that value.

Note: If the threshold is exceeded, all data points are displayed.

- Bounded Trigger Threshold
  - Create chart *only when there are data points outside a range you select.*

Example: You want to generate a certain chart only when X > 100 or X < 10. With this option selected, the chart is generated only when there are data points *outside* that range.

The four data points 9, 45, 67 and 99 would cause the chart to be generated (9 < 10). The data points 10, 45, 78 and 100 would *not* generate the chart (none of the data points is outside the range).

- Display Threshold
  - Use to customize the threshold above which data are displayed.

Example: Choose this option to set a baseline.

Unlike the Trigger Threshold option, however, only data points above the configured threshold are displayed.

Analyzer: Chart (Basic)

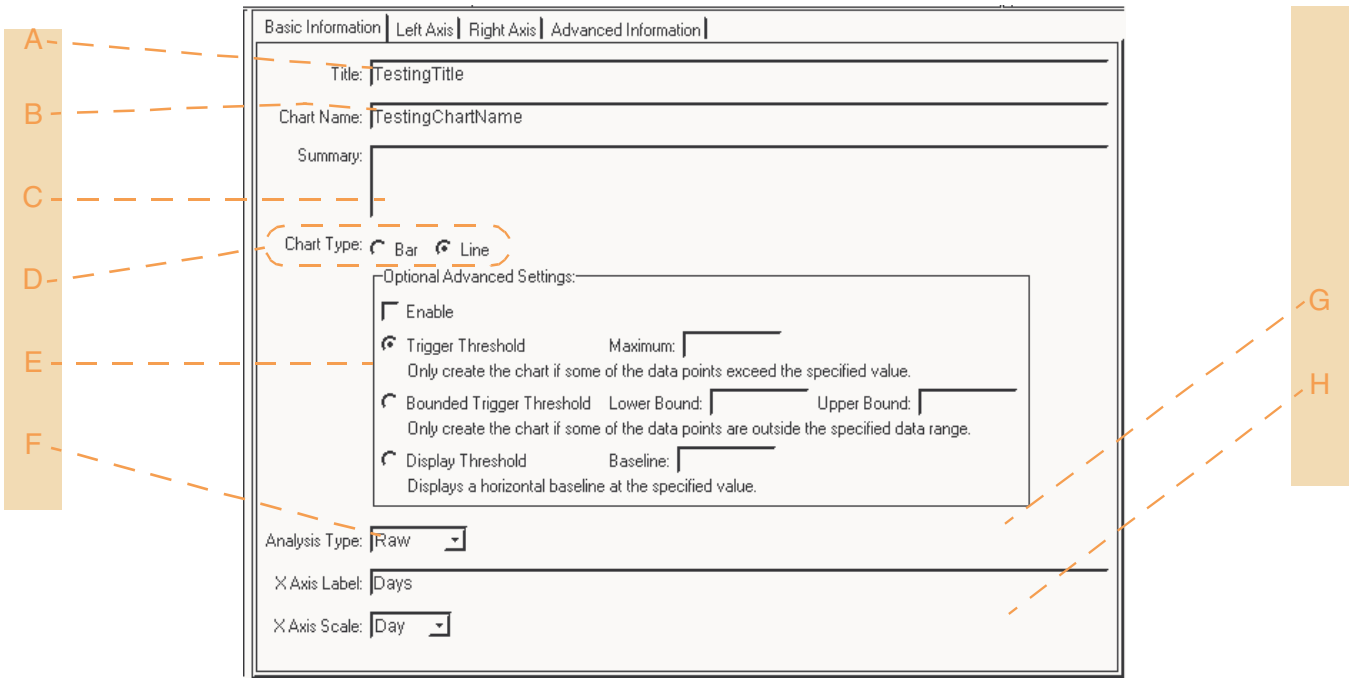

A: chart title
B: chart name
C: chart summary
D: use radio buttons to select chart type
E: if enabled (checked), one of three custom options can be selected here (see usage information on previous page, for details)
F: use combo box to select type of data (raw or average)
G: label for x (or horizontal) axis
H: time units applicable to x-axis (double-click icon for options)

Analyzer Dialogs: Chart (Left Axis)

Basics ...

Use this dialog to:

- enter a text label for the left-hand y-axis
- select the statistic(s) to associate with the left-hand y-axis
- set number of decimal places for displayed values
- select type of data to be displayed
- select units of value appropriate to the Range Type

How many statistics may be included:

If defined for a single server, up to 16 statistics may be associated with an axis. If any statistic is collected for multiple servers, only that one statistic may be associated with either the left or the right axis.

Putting it into practice ...

Range Type:

- *General* displays value in:
  - Thousandths
  - Ones
  - Thousands
  - Millions
- *Bytes* displays value in:
  - Bytes
  - KB
  - MB
  - GB
- Percent
  - units in % of applicable value(s)

(Range combo box not displayed if Range Type is Percent.)

Labels and the number of statistics reported:

The height of the graph (and the Y-axis labels to either side of it) is inversely proportional to the height of the legend. Eventually, the legend can grow so large, and the graph so small, that labels become truncated.

With line graphs, this situation is influenced both by the number of legend items and their length. With bar charts, where legend items are displayed vertically, it is the length of server/statistic names that causes the legend area to grow (one item with too long a name could cause label truncation).

Analyzer: Chart (Left Statistics)

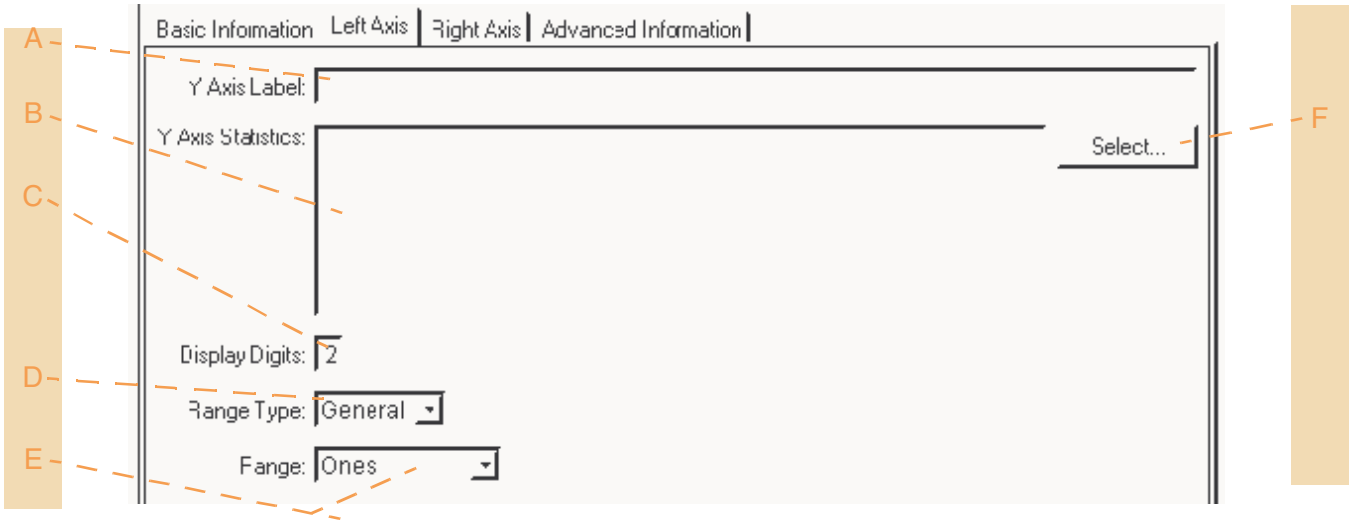

A: text label for the left-hand y-axis
B: statistics displayed by the left-hand y-axis
C: number of decimal places to display
D: type of data returned (general, bytes, or percent). For usage instructions, see previous page.
E: unit of value appropriate to Range Type. (For usage instructions, see previous page.)
F: button launches Select Statistic dialog

Analyzer Dialogs: Chart (Right Axis)

Basics ...

Use this dialog to:

- enter a text label for the right-hand y-axis
- select the statistic(s) to associate with the right-hand y-axis
- set number of decimal places for displayed values
- select type of data to be displayed
- select units of value appropriate to the Range Type

Putting it into practice ...

Range Type:

- *General* displays value in:
  - Thousandths
  - Ones
  - Thousands
  - Millions
- *Bytes* displays value in:
  - Bytes
  - KB
  - MB
  - GB
- Percent
  - units in % of applicable value(s)

The Range combo box is not displayed when Percent is selected as the Range Type.The Range combo box is not displayed when Percent is selected as the Range Type.

Maximum number of statistics per chart:

See *"How many statistics may be included:" on page 312*.

Potential label display issue:

See *"Labels and the number of statistics reported:" on page 312*.

Analyzer: Chart (Right Statistics)

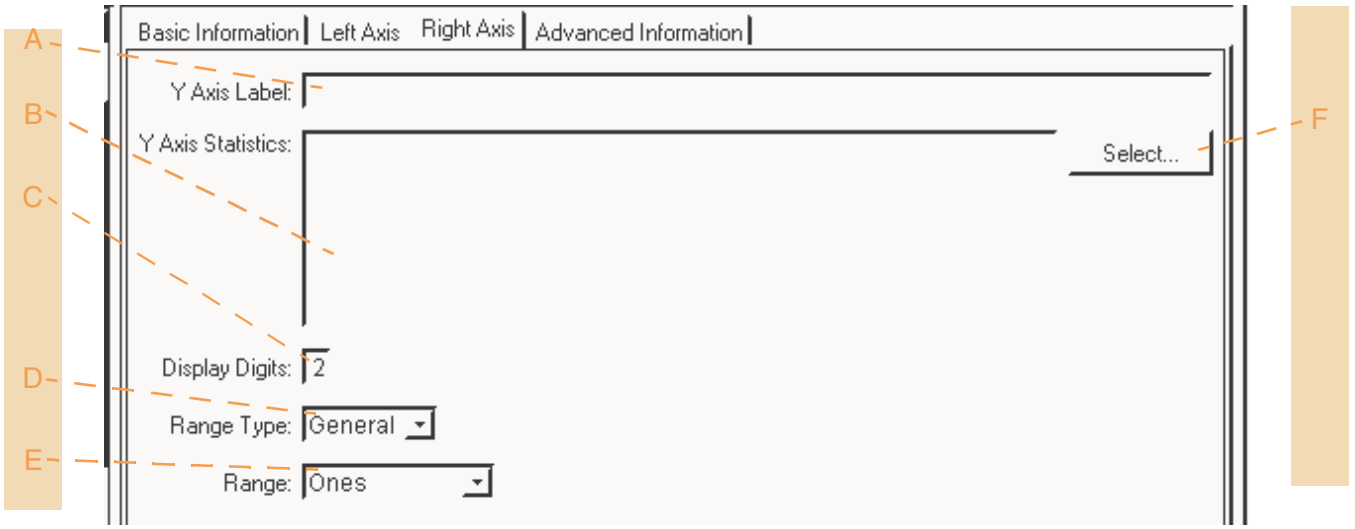

A: text label for the right-hand y-axis
B: statistics displayed by the right-hand y-axis
C: number of decimal places to display
D: type of data returned (see usage instructions, above)
E: unit of value appropriate to Range Type
F: button launches Select Statistic dialog

Analyzer Dialogs: Chart (Advanced Information)

Basics ...

Use to select options:

- save data retrieved as comma-delimited file
- custom day and date range for chart

How Custom date and time settings work together:

The date and time-of-day parameters under Custom provide two levels of filtering.

If you select 5/1/01 and 5/8/01 as the date bounds, and 8:00 AM and 5:00 PM as the time limits, for example, what does Analyzer return?

What it does ***not*** return is data from 8:00 AM on May 1st through 5:00 PM on May 8th.

Instead, for each day of the selected date range, Analyzer retrieves the requested data when it bears a time-stamp between the specified times, in this example between the hours of 8:00 AM and 5:00 PM.

Data bearing a time-stamp of 6:00 PM on May 7th, for instance, would not be retrieved.

Putting it into practice ...

Create CSV files:

Select if data retrieved to be imported into applications (such as Microsoft Excel) as comma-delimited lists.

Period Covered:

Broadly speaking, three options are available:

- Use Report Settings
  - Chart uses the date range of the Report.

    Example: When created, a chart was set up to return data for the previous month.

    Now, however, you want to incorporate this chart in a report covering only the current week.

    This can be accomplished by editing the chart, and setting Period Covered to *Use Report Settings*.

    Use this setting when creating charts that, by default, use the date range of the report in which it's included.
- Custom
  - Use dialog at B2 (on following page) to enter dates and times.
  - Custom *must* be used for date ranges longer than a month.
- Other periods (for list, see graphic B1 on next page)
  - Use the drop down in dialog (see graphic B3 on following page) to specify the range, and manually enter applicable custom times.

Analyzer: Chart (Advanced)

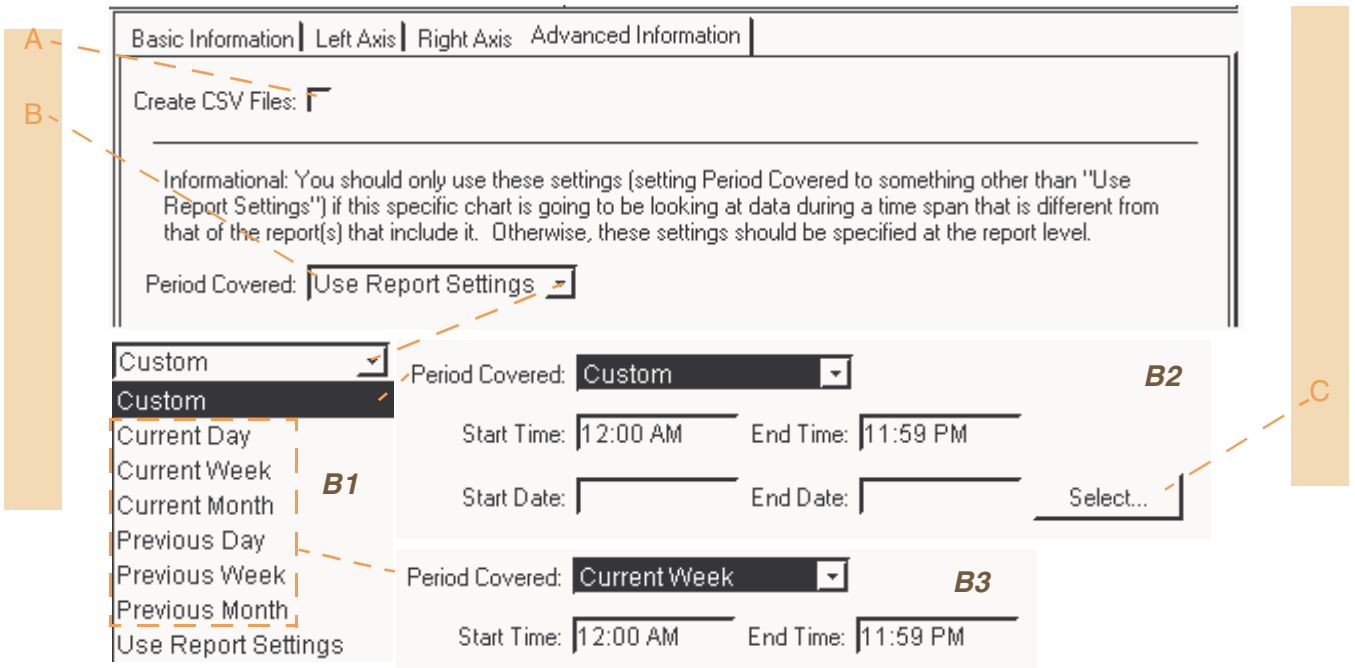

A: check to save data retrieved as comma-delimited file

B: use combo box to select period covered

If period =Use Report Settings, no date/time fields are displayed; if period=Custom, the fields on graphic ***B2*** are displayed; for all other periods, the fields on graphic ***B3*** are displayed.

C: button launches data-input dialog

Analyzer Dialogs: Report (Basic Information)

Basics ...

(For details on the following items, see ***Putting it into practice, below***.)

Use this dialog to enter:

- title of report
- name of .html file created
- report summary
- scheduling information
- mail recipients (by default, only the Admin receives a copy of the report)

Putting it into practice ...

Title vs Name fields:

The Title appears in two places:

- in the tree view, listing created reports
- at the top of the report when created as an .html file

The Report Name (plus the .html extension) is the file created and displayed.

Scheduling options:

The Schedule option selected influences what additional configuration options are displayed.

- Daily
  - No additional configuration fields.
- Weekly
  - Combo box to select the day of the week.
- nth Day of month
  - Text field in which a comma-delimited list of days (as integers) must be entered.
- One time only
  - No additional configuration fields.
- Do not run
  - No additional configuration fields.

Analyzer: Report (Basic Information)

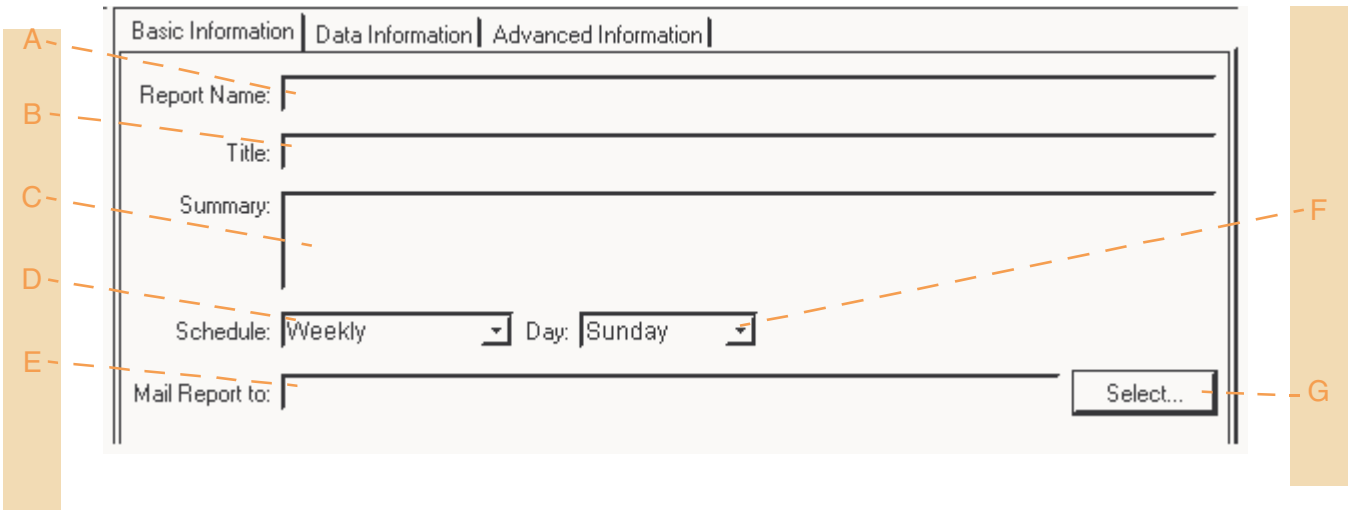

A: report name
B: report title
C: report summary
D: how often to run report (double-click icon for options)
E: comma-delimited list of mail recipients
F: day of the week on which to run report (applies to Weekly only)
G: button launches Select Mail Addresses dialog

Analyzer Dialogs: Report (Data Information)

Basics ...

Supplies report with:

- list of charts to be included
- date and time settings at report level

Putting it into practice ...

Selecting charts:

Chart selection must be done via the dialog (launched by clicking the Select... button.

Chart Titles cannot be filled in by hand.

Override Chart Settings:

The date/time settings used by reports depend, in part, on the settings specified at the chart level. The possible combinations are set forth in the table, below.

Please note that, regardless of the settings in the charts it incorporates, a report cannot be saved without a specified *Period Covered*, and both a Start and an End Time. By this means, reports always have default settings to which they can revert, if chart-level settings do not take precedence.

Similarly, charts where the *Period Covered* is not "Use Report Settings' cannot be saved if any date/time fields are left blank.

Report vs Chart settings

Report level	Chart level	Settings used by report
"Override..." selected	Any setting	Report-level settings
"Override..." deselected	Use Report Settings	Report-level settings
"Override..." deselected	Any specified *Period Covered* (including all associated settings)	Chart-level settings

Analyzer: Report (Data Information)

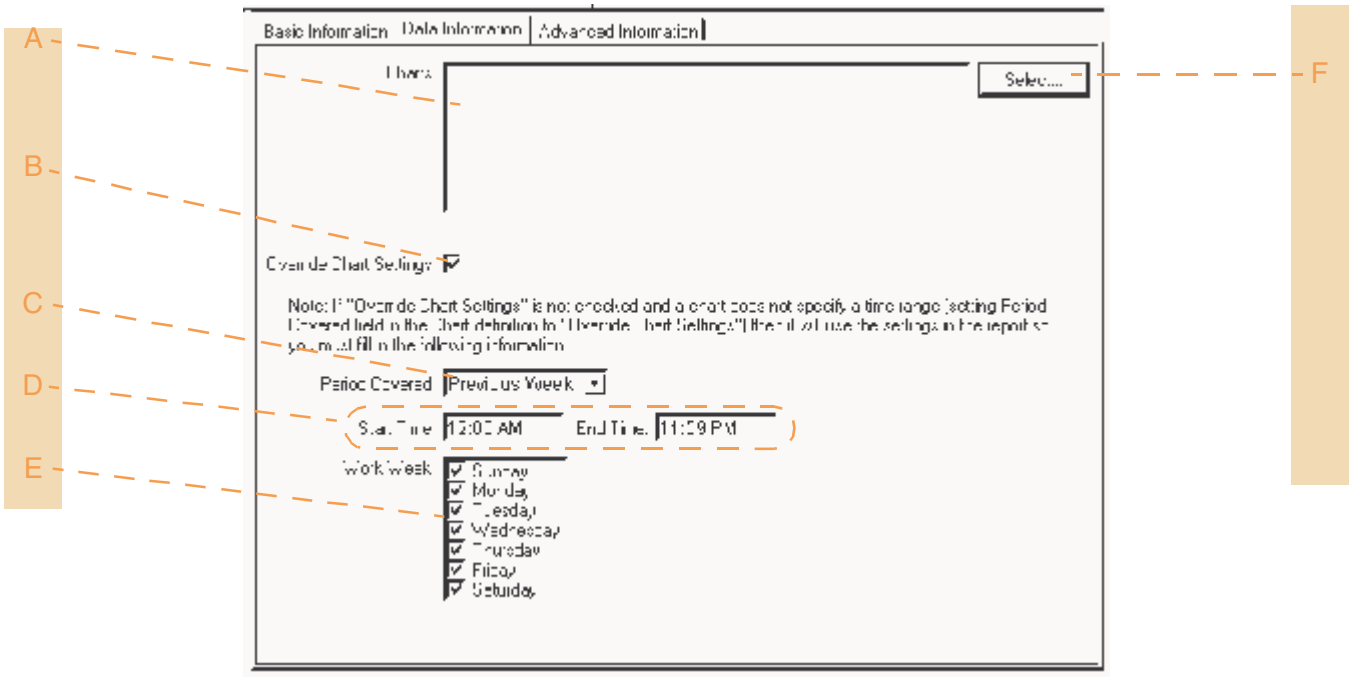

A: charts to be included in report
B: check to override chart settings for day/date and time ranges
C: period covered at report level
D: start/end times for report data to be included
E: select the days that constitute your work week
F: button launches Select Charts dialog

Analyzer Dialogs: Report (Advanced Information)

Basics ...

Use to enter custom settings re:

- creation of comma-delimited files for retrieved data
- location of statrep.nsf
- location of iwstats.nsf
- location for storing Analyzer reports

CSV file size:

The size of CSV files can grow quite large (10s of MBs), if charts contain large numbers of datapoints. For each datapoint, a separate record is created, including:

- StatisticDate
- StatisticName (255-character maximum)
- StatisticValue
- ServerName (255-character maximum)

A chart with 50,000 datapoints, saved as a CSV file and imported into Excel, for instance, would generate an *.xls file with 50,000 rows.

Putting it into practice ...

CSV options:

- Yes
  - overrides chart-level settings

Comma-delimited files are created for ALL charts for which data are (successfully) retrieved.

- No
  - overrides chart-level settings

NO comma-delimited files are created, even if, at the chart level CSV-file creation is selected.

- Chart
  - CSV-file creation determined by chart-level settings.

Example: Report ABC includes Charts 1, 2, and 3. At the chart level, CSV-file creation is turned on for Chart 3 only.

When Report ABC is run, a CSV file is created for Chart 3, but not for Charts 1 and 2.

Non-standard locations:

Make entries in fields B-D on the dialog (see following page) *only* if the item deviates from the default name/location.

Analyzer: Report (Advanced Information)

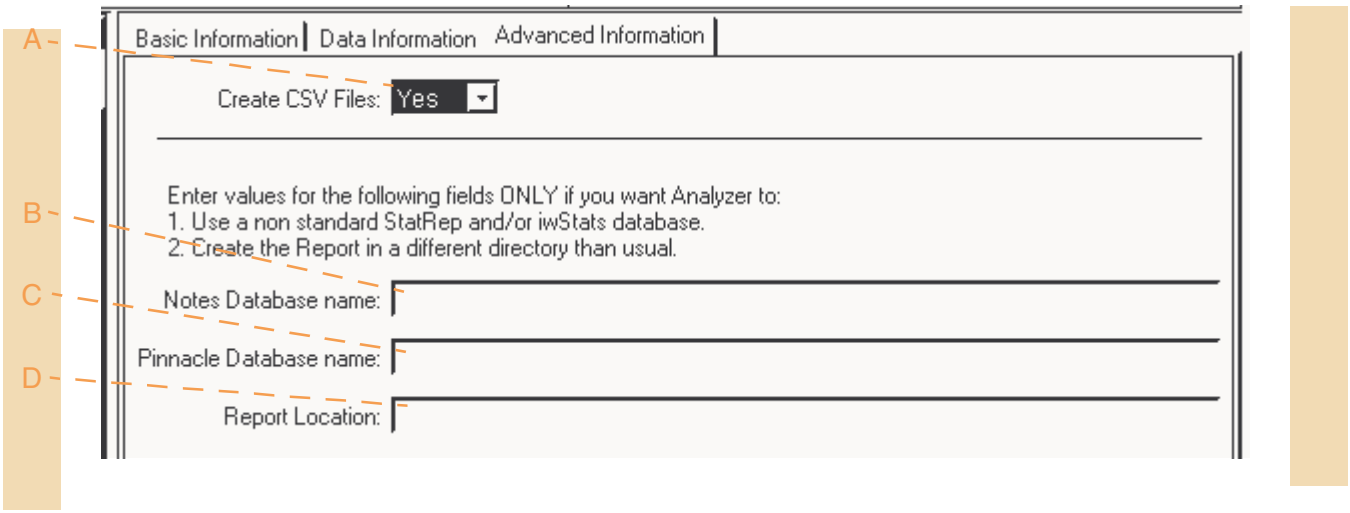

A: selection controls creation of comma-delimited files of retrieved data (yes, no, chart)

B: enter a database ONLY when it differs from the default database (and/or location)

C: enter a database ONLY when it differs from the default database (and/or location)

D: enter a location ONLY when it differs from the default

Chapter

10 Configuring IntelliWatch via Notes

Most Notes administrators use the Pinnacle Console to configure IntelliWatch components. A few, however, prefer to use a Notes client instead. For this reason, a summary is provided here on configuring IntelliWatch via Notes.

Chapter Contents

10.1.0.0 OVERVIEW

While not all operations discussed in the preceding chapters can be done through a Notes client, the following procedures can be performed, and are categorized by component:

- Management Agents
  - Triggers

    Editing Triggers
    Enabling/Disabling Triggers
    Creating Triggers
    Copying Triggers
    Deleting Triggers
  - Commands

    Editing Commands
    Creating Commands
    Copying Commands
    Deleting Commands
- ASW
  - Adding ASW Hubs
  - Deleting ASW Hubs
  - Activate/Deactivate Servers to Monitor
  - Maintenance Profiles

    Creating Maintenance Profiles
    Editing Maintenance Profiles
  - Action Profiles

    Creating Action Profiles
    Editing Action Profiles
- Pinnacle Performance Manager
- Analyzer
- Loading products at the admin console

10.2.0.0 MANAGEMENT AGENTS

10.2.1.0 Triggers

10.2.1.1 Editing Triggers

Editing Triggers is usually a straightforward procedure, involving modifications to a few fields (or perhaps only the server list, or a single parameter).

TO EDIT TRIGGERS:

1. Open the database containing the Trigger to be edited by going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > [Filename of MA]** via the drop-down menus.
2. Drill down by category/type and highlight the Trigger you want to edit.

Try searching via the Enabled Triggers view, or create your own Trigger categories, with Descriptions that simplify Trigger searches.

Remember ...

The Description field (when accessing MAs via a Notes client) is equivalent to the Names field (when accessing MAs via the Pinnacle Console).

3 Click on the 

toolbar button; alternatively, go to **Actions > Edit Trigger** via the drop-down menus.

4 Make any necessary changes to the Trigger fields.

5 Save the document by doing one of the following:

- press Escape
  - A dialog is displayed asking you if you want to save your changes (Yes is the default). The document then closes without further actions.
- click on the 

toolbar button
- go to **Actions > Save** via the drop-down menus
  - Saving documents in either of the last two ways requires you subsequently to close them, either by pressing Escape, by going to **File > Close** via the drop-down menus, or by clicking on the 

toolbar icon.

10.2.1.2 Enabling/Disabling Triggers

The importance of managing which Triggers on your system are enabled/disabled should not be underestimated.

As with any application/process, Trigger evaluation and execution takes up system resources. To make the most efficient use of Monitor Triggers, spend a few extra minutes when selecting the servers for which Triggers are to be enabled/disabled.

TO ENABLE/DISABLE TRIGGERS:

1 Open the database containing the Trigger(s) to be enabled/disabled by going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > [Filename of MA]** via the drop-down menus.

2 Drill down by category/type and highlight the Trigger(s) you want to enable/disable.

To disable Triggers, search using the Enabled Triggers view.

To enable a Trigger or Triggers, search using the Disabled Triggers view.

3 Enable/disable highlighted Triggers in one of the following ways:

- click on the 

(or Disable) toolbar button.
- go to **Actions > Enable** (or Disable) **Trigger** via the drop-down menus.

If you've selected more than one Trigger, use the menu option **Actions > Enable** (or Disable) **Selected Triggers**.

- Open the Trigger document and select the Enabled (or Disabled) radio button in the Basic Trigger Information section, then confirm the action in one of the following ways:
- press Escape
  - A dialog is displayed asking you if you want to save your changes (Yes is the default). The document then closes without further actions.

- click on the  toolbar button
- go to **Actions > Save** via the drop-down menus

  – Saving documents in either of the last two ways requires you subsequently to close them, either by pressing Escape, by going to **File > Close** via the drop-down menus, or by clicking on the  toolbar icon.

10.2.1.3 Creating Triggers

The first step is deciding in which database you want to create a Trigger. The primary factor is how *you* choose to organize your Triggers.

Although the seven default MAs are organized by areas of Domino functionality, *you can create any Trigger type in any MA.* In other words, you can create a Trigger governing database corruption in the Core Server MA, even though there is a separate Database Corruption MA.

TO CREATE TRIGGERS:

1. Open the database in which the Trigger is to be created by going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > [Filename of MA]** via the drop-down menus.
2. At the menu bar, select **Create > Trigger > [Desired Trigger Type, or Other for more types]**.
3. Fill in fields using guidelines set out in *"Chapter 3, Management Agents"*, specifically the sections on *"Creating Triggers" on page 54* and *"Trigger Condition fields" on page 56*.
4. Save the Trigger in one of the following ways:

- press Escape

  – A dialog is displayed asking you if you want to save your changes (Yes is the default). The document then closes without further actions.

- click on the  toolbar button
- go to **Actions > Save** via the drop-down menus

  – Saving documents in either of the last two ways requires you subsequently to close them, either by pressing Escape, by going to **File > Close** via the drop-down menus, or by clicking on the  toolbar icon.

10.2.1.4 Copying Triggers

Perhaps the simplest way of creating new Triggers is to copy pre-existing ones, whether within the same MA, or into a different/new MA.

Copying Triggers is especially useful when the new Trigger differs from an existing one in only one or two parameters.

to copy triggers:

1. Open the database containing the Trigger(s) to be copied by going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > [Filename of MA]** via the drop-down menus.
2. Select the Trigger(s) to be copied from the list.
3. Select **Edit > Copy** via the drop-down menus (or use Ctrl-c).
4. Switch to the new MA database.

5 Select **Edit > Paste** (or use Ctrl-v).

As soon as you finish copying the Trigger, change the Name in a way that makes the purpose of the Trigger clear.

Remember ...

If you forget to change the Name, no error occurs, but the Trigger's Name and function are not in agreement.

10.2.1.5 Deleting Triggers

We suggest disabling Triggers rather than deleting them, *since there is no Undo for Trigger deletion*.

Deleting a Trigger affects only the open MA. All other MAs that contained the Trigger, including the MA template, still contain it. Moreover, subject to how you manage your MAs, you can restore a deleted Trigger via Notes replication.

TO DELETE TRIGGERS:

1 Open the database containing the Trigger(s) to be deleted by going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > [Filename of MA]** via the drop-down menus.

2 Drill down by Trigger category/type and highlight the Trigger(s) you want to delete.

3 Do one of the following to initiate deletion:

- go to **Actions > Delete Trigger** via the drop-down menus
- click on the toolbar button

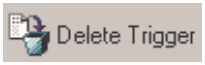

4 Proceed by doing one of the following:

- press **F9**
- go to **View > Refresh** via the drop-down menus
  - A confirmation dialog (default Yes) is displayed *before* the View is closed.
- press Escape
  - A confirmation dialog (default Yes) is displayed *after* the View is closed.

5 Confirm deletion by clicking Yes (or cancel deletion by clicking No).

10.2.2.0 Commands

10.2.2.1 Editing Commands

Editing Commands is usually a straightforward procedure, often involving modification of a single field.

TO EDIT COMMANDS:

1 Open the Command database by going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > iwcmd.nsf** via the drop-down menus.

2 Drill down by type, and highlight the Command you want to edit.

3 Click on the 

toolbar button; alternatively, go to **Actions > Edit Command** via the drop-down menus.

Command button labels include the type (here, IWSleep).

4 Make any necessary changes to the Command fields.

5 Save the Command in one of the following ways:

- press Escape
  - A dialog is displayed asking you if you want to save your changes (Yes is the default). The document then closes without further actions.
- click on the  toolbar button
- go to **Actions > Save** via the drop-down menus
  - Saving documents in either of the last two ways requires you subsequently to close them, either by pressing Escape, by going to **File > Close** via the drop-down menus, or by clicking on the  toolbar icon.

10.2.2.2 Creating Commands

Before creating a new Command, give some thought to a Name that is *short* and *unique*.

The Name field is used to populate the Available Commands list box; the more descriptive the Name, the easier it is to select Commands quickly and accurately.

TO CREATE COMMANDS:

1 Open the Command database by going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > iwcmd.nsf** via the drop-down menus.

2 At the menu bar, select **Create > Command >** [**Desired Command Type,** or **Other** for more types].

3 Fill in fields using guidelines set out in *"Creating Commands" on page 58*.

4 Save the Command in one of the following ways:

- press Escape
  - A dialog is displayed asking you if you want to save your changes (Yes is the default). The document then closes without further actions.
- click on the  toolbar button
- go to **Actions > Save** via the drop-down menus
  - Saving documents in either of the last two ways requires you subsequently to close them, either by pressing Escape, by going to **File > Close** via the drop-down menus, or by clicking on the  toolbar icon.

10.2.2.3 Copying Commands

Perhaps the simplest way of creating new Commands is to copy pre-existing ones.

Commands have many fewer configurable parameters, so copying them as opposed to creating them saves you less time than is the case with Triggers. Nevertheless, there may be situations when you want to copy rather than create Commands.

to copy commands:

1 Open the Command database by going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > iwcmd.nsf** via the drop-down menus.

2 Select the Commands(s) to be copied from the list.
3 Select **Edit > Copy** via the drop-down menus (or use Ctrl-c).
4 Select **Edit > Paste** (or use Ctrl-v).
5 Double-click the copy to open it, then click on the 

toolbar button; alternatively, go to **Actions > Edit [Type] Command** via the drop-down menus.

Command buttons and the relevant Action menu items the type (here, IWSleep).

6 Change the Name, and any other parameters, as per the purpose of the new Command.
7 Save the Command in one of the following ways:

- press Escape
  - A dialog is displayed asking you if you want to save your changes (Yes is the default). The document then closes without further actions.
- click on the 

toolbar button
- go to **Actions > Save** via the drop-down menus
  - Saving documents in either of the last two ways requires you subsequently to close them, either by pressing Escape, by going to **File > Close** via the drop-down menus, or by clicking on the 

toolbar icon.

10.2.2.4 Deleting Commands

We suggest deselecting Commands rather than deleting them. For three important reasons:

Deleting a Command affects all local MAs whose Triggers invoke it.

Deleted Commands cannot be restored.

Until you replicate **iwcmd.nsf**, a deleted Command still exists in copies of the database residing on other servers. Subject to how you manage your IntelliWatch databases, you can restore the Command via Notes replication.

TO DELETE COMMANDS:

1 Open the Command database by going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > iwcmd.nsf** via the drop-down menus.
2 In the left-hand pane, select the View corresponding to the type of the Command you want to delete.

If only the Commands folder is visible, click on it to see Views by Command type.

3 In the right-hand pane, select the Command(s) you want to delete.
4 Do one of the following to initiate deletion:

- go to **Actions > Delete Command** via the drop-down menus
- click on the toolbar button 

  - The label of the Delete button includes a reference to the Command type (here Restart Add-in).

5 Proceed by doing *one* of the following:

- press **F9**
- go to **View > Refresh** via the drop-down menus
  - A confirmation dialog (default Yes) is displayed *before* the View is closed.

- press Escape

  – A confirmation dialog (default Yes) is displayed *after* the View is closed.

6 Confirm deletion by clicking Yes (or cancel deletion by clicking No).

10.2.3.0 Creating New MAs via Notes Client

IntelliWatch Monitor provides predefined MAs that should meet the needs of the majority of Notes administrators. We recognize, however, that your Notes management needs may best be served by creating additional MAs.

Using the templates provided, creating new MAs is quick and easy.

TO CREATE NEW MAS:

1 From the Notes Console, select **File > Database > New**. This displays the New Database dialog.

2 Enter or select the following information at the dialog:

- server where custom MA resides

  If necessary, click button to expand list.
- title (Name) for custom MA
- filename for custom MA
- encryption and size limit settings, if desired
- to enhance Trigger searches, select *Create a full-text index for searching*
- in template list box, select *Intelliwatch MA Template*

  All Intelliwatch MAs are based on this template.
- select Template Server button only if template doesn't reside on server where MA is being created

If you are unfamiliar with these settings, please consult your Notes Client documentation.

3 Click Ok.

- opens newly created MA, and displays *About [MA Name]* screen

The new MA does not contain any Triggers! You must either create them in the new MA, or paste them in from another Management Agent.

As with the predefined MAs, all currently existing Commands (in iwcmd.nsf) are available in Triggers of the new MA (in the three *Then Take These Actions On The...* Occurrence sections).

You can now create Triggers in the new MA. However, IntelliWatch Monitor does not yet evaluate those Triggers. To make the new MA 'visible' to Monitor (and to the Pinnacle Console), you must take the following additional steps at the Pinnacle Console.

4 Click on the  toolbar button; alternatively, select **Solutions > Parameter Configuration** via the drop-down menus.

This displays the IntelliWatch Configuration Tool screen.

5 Select the name of the target system using the combo box at the upper left of the Console.

6 Click on the  toolbar icon.

7 In the left-hand pane, open the **Monitor** folder.
8 In the left-hand pane, select the value **IWAgent**.
9 In the right-hand pane, select the parameter **MA Database List** and click on the  toolbar icon (or double-click the parameter).
10 In the top field of the pop-up dialog, enter a comma-delimited list of all MAs whose Triggers you want Monitor to evaluate
11 In the bottom field of the dialog, type in a comma-delimited list of the servers on which you want the change made. Alternatively, use the Lookup button to launch a Select Servers dialog.

The MA Database List parameter will be modified in the iwparam.nsf of all servers in the list.

Iwagent refers to the MA Database List values to determine which Triggers should be evaluated. The Triggers in any MAs not in this list are not evaluated—even if they are enabled.

12 Click OK.
13 Stop and restart **iwagent** on the affected server(s) to read in the edited list of MAs immediately.

10.2.3.1 May I also create a custom Command database?

Yes. *However*, your new Command database must also be named **iwcmd.nsf**. Changing the name of the database causes a File Not Found error when the Pinnacle Console tries to connect to it.

10.3.0.0 ASW

All configuration of Advanced ServerWatch via a Notes client is done in the **iwasw.nsf** database, located in the same folder as the IntelliWatch MAs.

10.3.1.0 ASW Hubs

10.3.1.1 Adding ASW Hubs

There are two fundamental differences between setting up ASW Hubs via the Pinnacle Console versus via a Notes client:

- via the Pinnacle Console
  - – ASW Hubs created via the Pinnacle Console are displayed automatically.
  - – Until the Hub is assigned servers, no monitoring is taking place.

  For each server assigned to the Hub, a corresponding server document is created in that Hub's copy of **iwasw.nsf**.
- Via a Notes client
  - – ASW Hubs created via Notes are not automatically displayed at the Pinnacle Console. You must first do one of the following:
    - – Add them via the normal method at the Console.
    - – Add them to the **Pinnacle.ini** file under [ASW_GUI], using the syntax:

      Hub_Servers=[Server/Domain]~[Server/Domain]~[IP address]
  - – As soon as the first server document is created in a Hub's iwasw.nsf, monitoring is taking place.

The following procedure, "to add an ASW Hub", might therefore better be termed "initiating monitoring by ASW". Why is that? Because, broadly speaking, monitoring is

initiated only when the *first* server document is created in that Hub's **iwasw.nsf**. As soon as that database contains one server document, the ASW add-in task running on the system is monitoring. As additional server documents are created, all that changes is the *number* of servers being monitored.

TO ADD AN ASW HUB:

1 Open the Advanced ServerWatch database by either:

- double-clicking the workspace icon
- going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > iwasw.nsf.**

2 Open a server document in one of three ways:

- click the 

toolbar button
- go to **Create > Server** via the drop-down menus
- go to **Actions > Create Server** via the drop-down menus

3 Fill in *the first* of the following two fields:

- **Server**:
  - *Enter one server only*. Each server reporting to a given Hub requires its own document.
- Network Status (*cannot be edited*):
  - Populated and updated by the ASW Hub (at the user-specified ASW monitoring frequency).

4 Save the document, in one of two ways:

- press Escape and click the Yes button (the default)
- go to **File > Save** via the drop-down menus

5 Continue to add documents in this way, changing the name of the Monitoring Server whenever you want to create a new Hub.

New server documents become active at the next ASW monitoring cycle.

10.3.1.2 Deleting ASW Hubs

TO DELETE ASW HUBS:

1 Open the Advanced ServerWatch database by either:

- double-clicking the workspace icon
- going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > iwasw.nsf.**

2 Select *all* server documents in the ASW database on a given Hub, and delete them by going to **Edit > Cut** via the drop-down menus; alternatively, select the documents, then press **F9** (or **View > Refresh** via the drop-down menus), and confirm the deletion at the following dialog.

Remember ...

Changes must be made in the iwasw.nsf on the Hub server itself.

10.3.2.0 Managed Servers

10.3.2.1 Activate/Deactivate Servers to Monitor

With Pinnacle 99, each server document in the database contained a pair of radio

buttons that allowed users to activate/deactivate servers for monitoring.

This functionality has been simplified in IntelliWatch. If a server document exists in the Hub's **iwasw.nsf** database, that server is monitored by the local instance of ASW. If no server document is present, the server is not monitored.

TO ACTIVATE A SERVER FOR MONITORING:

1 Open the Advanced ServerWatch database by either:

- double-clicking the workspace icon
- going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > iwasw.nsf.**

2 Create a server document for the server the local instance of ASW is to monitor.

3 Save the document, in one of two ways:

- press Escape and click the Yes button (the default)
- go to **File > Save** via the drop-down menus

(Your changes take effect at the next ASW monitoring cycle.)

TO DEACTIVATE A SERVER FOR MONITORING:

1 Open the Advanced ServerWatch database by either:

- double-clicking the workspace icon
- going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > iwasw.nsf.**

2 Click on the Server List View in the left-hand pane.

3 Select the server document to be deleted, then:

- go to **Edit > Cut** via the drop-down menus

Be aware that this deletion method does not require confirmation.

- press the Delete button, followed by:
  - Escape, or
  - F9

  In either case, confirm the deletion by clicking the Yes button (the default)

10.3.3.0 Maintenance Profiles

This section discusses the creation and editing of Maintenance Profiles via a Notes client. For more details, see under *"Maintenance and Action Profiles" on page 176*.

- Creating Maintenance Profiles
- Editing Maintenance Profiles

10.3.3.1 Creating Maintenance Profiles

Maintenance Profiles are used to temporarily turn off ASW monitoring, so that a server down for maintenance does not generate alerts.

Maintenance Profiles offer four scheduling options: Daily, Weekly, Monthly, and Day of Month. Document fields vary slightly from one to the other.

- **Daily** is the most straightforward of the three, since it does not require the specification of a day of the week or month.
- Selecting **Weekly** brings up a collection of check boxes, one for each day of the week.
- **Monthly** brings up a field that takes a numerical value, based on the day of the month (that is, a *date*) on which to activate the Profile.

Enter it manually, or click on the arrow of the combo box to select from a list

- **Day of Month** has check boxes for both the days of the week and the weeks of the month.

  To activate a Maintenance Profile for the second and last Fridays of every month, select Friday in the upper group of checkboxes, and 2nd and Last in the lower group.

The following example assumes **Weekly** maintenance, to take place from 10:00-11:30pm on Fridays.

TO CREATE MAINTENANCE PROFILES:

1 Open the Advanced ServerWatch database by either:

- double-clicking the workspace icon, or
- going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > iwasw.nsf.**

2 Go to **Create > Maintenance Profile** via the drop-down menus. This opens a new Maintenance Profile document.

3 Enter a Profile Name that's *short, unique, and meaningful.*

4 Populate the Server Group field using the drop-down, or filling in groups/servers in by hand.

Multiple entries put in by hand should be comma-delimited.

5 Check Friday in the Day(s) of the Week field.

6 Enter 22:00 as the Down at time.

A 24-hour clock is assumed, unless you add AM or PM.

7 Enter 90 in the Duration field (value in minutes).

8 Save the document, in one of two ways:

- press Escape and click the Yes button (the default)
- go to **File > Save** via the drop-down menus

New Maintenance Profiles take effect as of the next ASW monitoring cycle, or when the scheduled Maintenance cycle begins, whichever is later.

10.3.3.2 Editing Maintenance Profiles

Once you set up the Maintenance Profiles required for your environment, most cases of adding servers, or of changing Maintenance parameters, can be handled by editing an existing Maintenance Profile, rather than creating a new one.

TO EDIT MAINTENANCE PROFILES:

1 Open the Advanced ServerWatch database by either:

- double-clicking the workspace icon, or
- going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > iwasw.nsf.**

2 In the left-hand pane, expand Profiles, then select Maintenance Profiles, to display available documents.

3 Double-click the Profile to be edited, or go to **File > Open** via the drop-down menus.

4 Make the desired changes.

5 Save the document, in one of two ways:

- press Escape and click the Yes button (the default)
- go to **File > Save** via the drop-down menus

10.3.4.0 Action Profiles

This section discusses the creation and editing of Action Profiles via a Notes client. For more details on Action Profiles, see under *"Maintenance and Action Profiles" on page 176*.

- Creating Action Profiles
- Editing Action Profiles

10.3.4.1 Creating Action Profiles

Action Profiles allow you to tailor Notification options to the needs of a given Hub or set of monitored servers. In addition, using Start Program Commands, other programs/batch files can be launched to allow ASW's response to be customized to fit the needs of the environment.

While most Action Profile fields are self-explanatory, the following fields bear a few additional comments:

- Number of retries

  The number of times you want ASW to re-attempt the connection before deciding the server is down. If retries is set to zero, ASW executes the selected actions upon the first failed monitoring attempt.

- Treat BUSY as responding

  When ASW pings a server, it receives one of four responses:

  – RESPONDING

  – NOT_RESPONDING

  – BUSY

  – **RESTRICTED**.

  The **Treat BUSY as responding** option concerns the last two of these: BUSY and RESTRICTED.

  If this option is set to **True**, BUSY and RESTRICTED *are treated the same as RESPONDING,* that is, *no actions are taken*.

  If the option is set to **False**, BUSY and RESTRICTED *are treated the same as NOT_RESPONDING,* that is, *actions are taken*.

- Start Program

  Selecting this option activates an additional field, where you enter the name of a Program or Batch File you want to run whenever actions are taken by ASW.

10.3.4.2 Practical example

Assumptions:

- employ all notification methods except NT Event Log
- have ASW make one additional attempt before taking actions
- do *not* Treat BUSY as Responding
- run batch file **custom.bat** whenever actions are taken
- send e-mail notification to one additional address, Notes_help@company.com

If e-mail notification is selected, a message is automatically sent to the Notes Admin specified during Setup.

TO CREATE ACTION PROFILES:

1 Open the Advanced ServerWatch database by either:

- double-clicking the workspace icon, or
- going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > iwasw.nsf.**

2 Go to **Create > Action Profile** via the drop-down menus.

This opens a new Action Profile document.

3 Enter a Profile Name that's *short, unique and meaningful.*

4 Populate the Server Group field using the drop-down, or filling in groups/servers in by hand.

Multiple entries put in by hand should be comma-delimited.

5 Set the Number of Retries to one (1).

6 Set the Treat BUSY as Responding option to False.

7 Select all Actions *except* NT Event Log.

8 Enter *Notes_help@company.com in the E-mail Address List.*

9 Enter **custom.bat** in the Program to Run field.

10 Save the document, in one of two ways:

- press Escape and click the Yes button (the default)
- go to **File > Save** via the drop-down menus

10.3.4.3 Editing Action Profiles

TO EDIT ACTION PROFILES:

1 Open the Advanced ServerWatch database by either:

- double-clicking the workspace icon, or
- going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > iwasw.nsf.**

2 In the left-hand pane, expand Profiles, then select Action Profiles, to display available documents.

3 Double-click the Profile to be edited, or go to **File > Open** via the drop-down menus.

4 Make the desired changes.

5 Save the document, in one of two ways:

- press Escape and click the Yes button (the default)
- go to **File > Save** via the drop-down menus

Newly created Action Profiles take effect as of the next ASW monitoring cycle. The same is true of changes made to existing profiles.

10.4.0.0 PM

10.4.1.0 Statistics

All configuration of Pinnacle Performance Manager via a Notes client is done in the

iwpmstat.nsf database, located in the same folder as the IntelliWatch MAs.

The following operations can be performed at via the **iwpmstat.nsf** database:

- Creating Statistic Definitions
- Editing Statistic Definitions
- Deleting Statistic Definitions

10.4.1.1 Creating Statistic Definitions

The configuration parameters for PM statistics can vary considerably by type, making it overly complicated to produce a set of procedures that covers all statistics. Nevertheless, there are basic steps common to all types, and they will be set out here, using as an example the *Mail.Domain* type.

TO CREATE A PM STATISTIC DEFINITION:

1 Open the IntelliWatch statistic database by going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > iwpmstat.nsf** via the drop-down menus.

2 Go to **Create > Statistic > Mail Domain** via the drop-down menus; alternatively, go to the Statistics by Type folder in the left-hand pane, and click on the Mail Domain View.

A toolbar button is now displayed that allows you to create a statistic of this type.

3 If necessary, expand the Basic Statistic Information section by clicking on the green arrow to the left.

4 Give the statistic a unique Name, ideally, combining concision with a clear indication of the statistic's purpose in your environment.

In this example, you might call it *Mail.All.Org* to generate data for all domains with the.*org* suffix.

Remember ...

The Name is what PM assigns to the generated data, and what Analyzer must subsequently use to retrieve it.

5 Give the statistic a description that complements the information provided by the Name.

6 If necessary, expand the Statistic Collection Information section by clicking on the green arrow to the left.

7 Select the servers/groups for which to generate data, or leave blank for all.

8 Select the *True* radio button to make generated data available to Analyzer, or *False* if data need not be collected for reporting purposes.

9 Select the *True* radio button to make generated data available for in-memory monitoring by Triggers, or *False* if the statistic data need not be available for that purpose.

In this example, you might want to use Analyzer to report on the data, but may not need to have the statistic monitored by IntelliWatch Triggers.

10 Select the radio button for the applicable Notes version(s).

11 If necessary, expand the Details section by clicking on the green arrow to the left.

12 Enter the *org* suffix for which this statistic is to generate data.

13 Save the document by doing one of the following:

- press Escape
  - – A dialog is displayed asking you if you want to save your changes (Yes is the default). The document then closes without further actions.

- click on the  toolbar button
- go to **Actions > Save** via the drop-down menus

  – Saving documents in either of the last two ways requires you subsequently to close them, either by pressing Escape, by going to **File > Close** via the drop-down menus, or by clicking on the  toolbar icon.

10.4.1.2 Editing Statistic Definitions

The statistic that will be edited in this example is *Mail.Incoming.Exceeds.OneHour* (of the *Mail.Incoming.Delivery* type).

You want to change the statistic's time interval from one hour to thirty minutes.

TO EDIT A PM STATISTIC DEFINITION:

1 Open the IntelliWatch statistic database by going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > iwpmstat.nsf** via the drop-down menus.

2 Go to the Statistics by Type folder in the left-hand pane, and click on the Mail Incoming Delivery View.

3 In the right-hand pane, select the statistic to be modified, and click on the toolbar button to put it into edit mode; alternatively, go to **Actions > Edit Statistic** via the drop-down menus.

Several steps from the Statistic Creation procedure are skipped here, since only two fields are going to be changed: the Name (to reflect the new time interval) and the interval itself.

4 If necessary, expand the Basic Statistic Information section by clicking on the green arrow to the left.

5 Modify the suffix of the statistic name to *.HalfHour*.

6 If necessary, expand the Details section by clicking on the green arrow to the left.

7 Change the *Time Over (minutes)* parameter from 60 to 30.

8 Save the document by doing one of the following:

- press Escape

  – A dialog is displayed asking you if you want to save your changes (Yes is the default). The document then closes without further actions.

- click on the  toolbar button
- go to **Actions > Save** via the drop-down menus

  – Saving documents in either of the last two ways requires you subsequently to close them, either by pressing Escape, by going to **File > Close** via the drop-down menus, or by clicking on the  toolbar icon.

10.4.1.3 Deleting Statistic Definitions

All deletions follow the same steps. The statistic that will be deleted in this example is of the Difference type.

TO DELETE A PM STATISTIC DEFINITION:

1 Open the IntelliWatch statistic database by going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > iwpmstat.nsf** via the drop-down menus.

2 Go to the Statistics by Type folder in the left-hand pane, and click on the Difference View.

3 In the right-hand pane, select the statistic(s) to be deleted, then do one of the following:

- click on the Delete Statistic toolbar icon
- go to **Actions > Delete Statistic** via the drop-down menus
- press the Delete key

The statistic is now marked for deletion.

4 Complete the deletion procedure by doing one of the following:

- press Escape
- press F9
- exit the Notes client

  – In all cases, a dialog is displayed asking you if you want to delete the selected statistic(s). Clicking on Yes (the default) completes the deletion.

10.4.1.4 Configuration details by type

More information on configuring individual statistic types is available in *"Chapter 5, Performance Manager" on pages 218 to 249*.

10.5.0.0 ANALYZER

10.5.1.0 Getting started

As discussed in *"Chapter 9, Analyzer"*, effective report creation begins at the statistic level.

Simply put, chart and report generation cannot succeed unless *all* of the following conditions are fulfilled:

- location of data

  – Statistics to be reported on must reside on the system running Analyzer Server (in **iwstats.nsf** and **statrep.nsf**). IntelliWatch's Analyzer does not retrieve statistics from databases on other systems.

- statistic Names/Titles

  – Analyzer statistic Names must correspond *exactly* with the PM/Notes statistic they are intended to retrieve.

  – Statistic Titles (specific to Analyzer), on the other hand, need *not* correspond to the PM name.

  MyStat.Test could serve as the Title for the statistic named *Disk.D.Free*, for example. As long as the statistic collected using PM also has the Name *Disk.D.Free*, data retrieval will succeed (assuming all other conditions are met).

- source server(s)

  – The relevant statistics must have been collected for the specified server(s).

  If data is present for Srvr1, but not Srvr2, chart creation will fail if only the latter server is specified.

- date/time

  – The relevant statistics must have been collected for the specified date(s)/time(s).

  If statistics have been collected for the specified servers for the Current Week, but not the Previous Week, chart creation will fail if the latter period is specified.

10.5.2.0 Statistics

The following operations can be performed via the **analyzer.nsf** database:

- Creating Statistic Definitions
- Editing Statistic Definitions
- Deleting Statistic Definitions

10.5.2.1 Creating Statistic Definitions

All Analyzer statistic definitions are created using the same procedure.

TO CREATE AN ANALYZER STATISTIC DEFINITION:

1 Open the Analyzer database by going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > analyzer.nsf** via the drop-down menus.

2 Go to **Create > Statistic** via the drop-down menus.

3 Give the statistic a unique Title, ideally, combining concision with a clear indication of the statistic's purpose in your environment.

A Title such as *AllSvrs.DskC.Free* might be useful, for example.

Remember ...

The Title need not correspond in any way to the statistic Name in either Analyzer or PM.

4 Enter the Name of the Pinnacle/Notes statistic to associate with this Analyzer definition.

5 Select the servers/groups for which to retrieve data, or leave blank for all.

6 Using the combo box, select the source database.

Pinnacle statistic=**iwstats.nsf**; Notes statistic=**statrep.nsf**.

7 Save the statistic definition by doing one of the following:

- press Escape
  - A dialog is displayed asking you if you want to save your changes (Yes is the default). The document then closes without further actions.
- click on the 

toolbar button
- go to **Actions > Save** via the drop-down menus
  - Saving documents in either of the last two ways requires you subsequently to close them, either by pressing Escape, by going to **File > Close** via the drop-down menus, or by clicking on the 

toolbar icon.

10.5.2.2 Editing Statistic Definitions

Editing procedures are identical for all Analyzer statistic definitions.

TO EDIT AN ANALYZER STATISTIC DEFINITION:

1 Open the Analyzer database by going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > analyzer.nsf** via the drop-down menus.

2 In the left-hand pane, click on the Statistics View.

3 In the right-hand pane, select the statistic definition to be edited.

4 Put the statistic definition into edit mode by doing one of the following:

- double-click the statistic, then click on the Edit toolbar button
- go to **Actions > Edit Document** via the drop-down menus

5 Make the desired modifications to the statistic definition.

6 Save the statistic definition by doing one of the following:

- press Escape
  - A dialog is displayed asking you if you want to save your changes (Yes is the default). The document then closes without further actions.
- click on the 

toolbar button
- go to **Actions > Save** via the drop-down menus
  - Saving documents in either of the last two ways requires you subsequently to close them, either by pressing Escape, by going to **File > Close** via the drop-down menus, or by clicking on the 

toolbar icon.

10.5.2.3 Deleting Statistic Definitions

All deletions of Analyzer statistic definitions follow the same steps.

TO DELETE AN ANALYZER STATISTIC DEFINITION:

1 Open the Analyzer database by going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > analyzer.nsf** via the drop-down menus.

2 In the left-hand pane, click on the Statistics View.

3 In the right-hand pane, select the statistic(s) to be deleted, then press the Delete key.

The statistic is now marked for deletion.

4 Complete the deletion procedure by doing one of the following:

- press Escape
- press F9
- exit the Notes client
  - In all cases, a dialog is displayed asking you if you want to delete the selected statistic(s). Clicking on Yes (the default) completes the deletion.

10.5.2.4 Configuration details by type

More information on configuring individual statistic definitions is available in *"Chapter 9, Analyzer"* under *"Defining Statistics" on page 299*.

10.5.3.0 Charts

- Creating Charts
- Editing Charts
- Deleting Charts

10.5.3.1 Creating Charts

All Analyzer charts are created using the same procedure.

TO CREATE AN ANALYZER CHART:

1 Open the Analyzer database by going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > analyzer.nsf** via the drop-down menus.

2 Go to **Create > Chart** via the drop-down menus.

3 Give the chart a unique Title, ideally, combining concision with a clear indication of the chart's purpose.

The Title appears in two places: 1) in the tree view, listing created charts; 2) at the top of the chart when created as a .gif file.

4 Enter the Name of the Pinnacle/Notes chart.

The Chart Name (plus the .gif extension) is the file created and displayed by the report.

5 Fill in the remaining chart fields. Refer, as necessary, to *"Chapter 9, Analyzer"*, specifically to section *9.4.2.0 on page 300*, as well the dialog-usage instructions *on pages 310 to 317*.

6 Save the chart by doing one of the following:

- press Escape
  - A dialog is displayed asking you if you want to save your changes (Yes is the default). The document then closes without further actions.
- click on the 

toolbar button
- go to **Actions > Save** via the drop-down menus
  - Saving documents in either of the last two ways requires you subsequently to close them, either by pressing Escape, by going to **File > Close** via the drop-down menus, or by clicking on the 

toolbar icon.

10.5.3.2 Editing Charts

Editing procedures are identical for all Analyzer charts.

TO EDIT AN ANALYZER CHART:

1 Open the Analyzer database by going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > analyzer.nsf** via the drop-down menus.

2 In the left-hand pane, click on the Charts View.

3 In the right-hand pane, select the chart to be edited.

4 Put the chart into edit mode by doing one of the following:

- double-click the chart, then click on the Edit toolbar button
- go to **Actions > Edit Document** via the drop-down menus

5 Make the desired modifications to the chart.

6 Save the chart by doing one of the following:

- press Escape
  - A dialog is displayed asking you if you want to save your changes (Yes is the default). The document then closes without further actions.
- click on the 

toolbar button
- go to **Actions > Save** via the drop-down menus
  - Saving documents in either of the last two ways requires you subsequently to close them, either by pressing Escape, by going to **File > Close** via the drop-down menus, or by clicking on the 

toolbar icon.

10.5.3.3 Deleting Charts

All deletions of Analyzer charts follow the same steps.

TO DELETE AN ANALYZER CHART:

1 Open the Analyzer database by going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > analyzer.nsf** via the drop-down menus.

2 In the left-hand pane, click on the Charts View.

3 In the right-hand pane, select the chart(s) to be deleted.
4 In the right-hand pane, select the chart(s) to be deleted, then press the Delete key.

The chart is now marked for deletion.

5 Complete the deletion procedure by doing one of the following:

- press Escape
- press F9
- exit the Notes client

  – In all cases, a dialog is displayed asking you if you want to delete the selected chart(s). Clicking on Yes (the default) completes the deletion.

10.5.4.0 Reports

10.5.4.1 Creating Reports

All Analyzer reports are created using the same procedure.

TO CREATE AN ANALYZER REPORT:

1 Open the Analyzer database by going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > analyzer.nsf** via the drop-down menus.
2 Go to **Create > Report** via the drop-down menus.
3 Give the report a unique Title, ideally, combining concision with a clear indication of the report's purpose.

The Title appears in two places: 1) in the tree view, listing created reports; 2) at the top of the report when created as an .html file

4 Enter the Name of the Pinnacle/Notes report.

The Report Name (plus the .html extension) is the file created and displayed.

5 Fill in the remaining report fields. Refer, as necessary, to *"Chapter 9, Analyzer"*, specifically to section *9.4.4.0 on page 303*, as well the dialog-usage instructions *on pages 318 to 323*.
6 Save the report by doing one of the following:

- press Escape

  – A dialog is displayed asking you if you want to save your changes (Yes is the default). The document then closes without further actions.

- click on the 

toolbar button
- go to **Actions > Save** via the drop-down menus

  – Saving documents in either of the last two ways requires you subsequently to close them, either by pressing Escape, by going to **File > Close** via the drop-down menus, or by clicking on the 

toolbar icon.

10.5.4.2 Editing Reports

Editing procedures are identical for all Analyzer reports.

TO EDIT AN ANALYZER REPORT:

1 Open the Analyzer database by going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > analyzer.nsf** via the drop-down menus.
2 In the left-hand pane, click on the Reports View.
3 In the right-hand pane, select the report to be edited.

4 Put the report into edit mode by doing one of the following:

- double-click the report, then click on the Edit toolbar button
- go to **Actions > Edit Document** via the drop-down menus

5 Make the desired modifications to the report.

6 Save the report by doing one of the following:

- press Escape

  – A dialog is displayed asking you if you want to save your changes (Yes is the default). The document then closes without further actions.

- click on the 

toolbar button
- go to **Actions > Save** via the drop-down menus

  – Saving documents in either of the last two ways requires you subsequently to close them, either by pressing Escape, by going to **File > Close** via the drop-down menus, or by clicking on the 

toolbar icon.

10.5.4.3 Deleting Reports

All deletions of Analyzer reports follow the same steps.

TO DELETE AN ANALYZER REPORT:

1 Open the Analyzer database by going to **File > Database > Open > [Name of Server where database resides] > [Name of IntelliWatch data directory] > analyzer.nsf** via the drop-down menus.

2 In the left-hand pane, click on the Reports View.

3 In the right-hand pane, select the report(s) to be deleted.

4 In the right-hand pane, select the report(s) to be deleted, then press the Delete key.

The report is now marked for deletion.

5 Complete the deletion procedure by doing one of the following:

- press Escape
- press F9
- exit the Notes client

  – In all cases, a dialog is displayed asking you if you want to delete the selected report(s). Clicking on Yes (the default) completes the deletion.

10.6.0.0 LOADING PRODUCTS AT THE ADMIN CONSOLE

The following procedure assumes you have all necessary rights to issue remote server commands on the machines concerned. If you follow these steps but encounter connection errors, please consult with the manager of your Notes team to obtain the required access privileges.

TO LOAD PRODUCTS:

1 Go to **File > Tools > Server Administration** via the drop-down menus.

2 Fill in the *Choose a server to administer* field.

3 Click on the *Console* button, which brings up a Remote Server Console.

4 Select *Live console* if you want to see server activity in real time.

5 Enter the desired Command in the *Server console command* text box.

Don't forget to use "load", "tell", and so on, as required by the Command. For example, to stop iwagent, type in tell iwagent q[uit].

6 Click the *Send* button.

Chapter

11 Console Utilities

Monitor Enterprise offers you several custom utilities that are accessed via the Pinnacle Console. These utilities do everything from remotely configuring your servers to sending SNMP traps.

Chapter Contents

11.1.0.0 PARAMETER CONFIGURATION UTILITY

11.1.1.0 Overview

On NT, local IntelliWatch settings are stored in the registry database.

On UNIX, these settings are found in a number of .ini files. Which files are present depends on the components installed. They include:

- iwasw.ini (Advanced ServerWatch)
- iwanlyzr.ini (Analyzer)
- iwmon.ini (Monitor)
- iwpinn.ini (general settings)
- iwperfmn.ini (Performance Manager)
- iwtracer.ini (Tracer)

The Parameter Configuration Utility provides administrators more convenient remote access to these settings.

Using the Parameter Configuration Utility instead of accessing settings directly in their repository offers administrators another advantage: *it's impossible to inadvertently change the configuration of other programs, since access is granted to IntelliWatch settings only.*

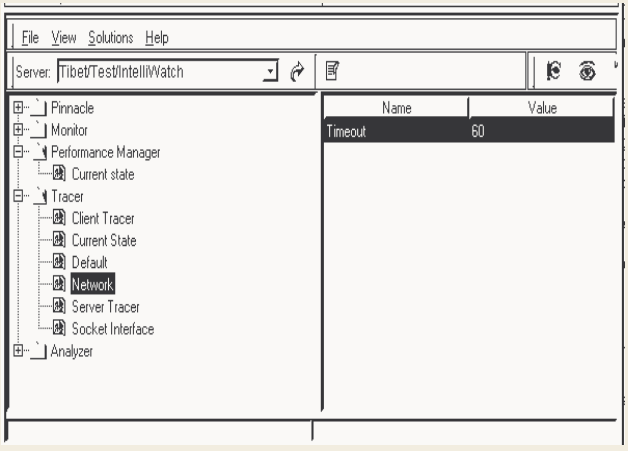

11.1.2.0 Using the Configuration Utility

1 Click on the Parameter Configuration toolbar icon.

This brings up the interface displayed below.

2 Use the combo box at the upper left to select the server.

3 With the correct server specified, click on the icon (to the right of the combo box).

This brings up a list of folders *for installed components*.

4 Access settings for components you want to configure by clicking on the + next to the folder. If there are sub folders, keep clicking the +'s until you see a list of current settings.

5 In the left-hand pane, select the parameter whose value is to be changed.

A value table is displayed in the right-hand pane (see the figure above, where the Network timeout parameter is currently selected).

6 In the right-hand pane, click on the parameter whose value is to be changed.

The background of the parameter changes from white to blue, and the edit icon is activated.

The parameter is not yet editable. To launch the Edit Parameter dialog, proceed with the following Step.

7 To change the parameter's value, click on the  icon; alternatively, double-click the parameter in the right-hand pane of the Console.

An Edit Parameter dialog is displayed, allowing you to enter a new value/values. For details, see *"Chapter 11, Edit Parameter dialogs"*, below.

8 Enter the new value(s) and click OK.

11.1.3.0 Edit Parameter dialogs

The dialog displayed differs according to the setting being changed.

For settings such as the boolean Show Server List[?], the dialog in *Figure 11-1* is displayed. A single value is required (here, true or false).

FIGURE 11-1: Edit Parameter dialog 1

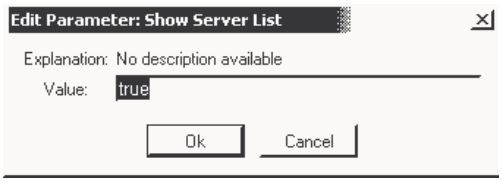

For a parameter like that governing the sender of Pinnacle e-mail messages, a dialog with two fields is displayed. Both the sender, and the server(s) on which that parameter should be set, are available for editing in *Figure 11-2*.

FIGURE 11-2: Edit Parameter dialog 2

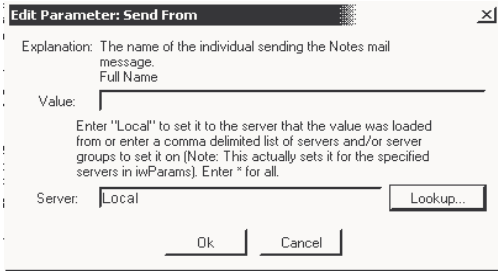

There are three options for the server field:

- Local
  - The *Mail* parameter is changed only in the **iwparam.nsf** of the server from which the value was loaded.
- Comma-delimited list
  - The *Mail* value is changed in the **iwparam.nsf** of all listed servers.
- Wildcard ('*')
  - The *Mail* value is changed in the **iwparam.nsf** of all known servers.

11.1.4.0 When Changes Take Effect

Configuration changes take effect the next time those settings are read by the relevant IntelliWatch component. If settings for Monitor (IWAGENT) or Advanced ServerWatch (IWASW) Add-in tasks are changed, they must be stopped and restarted before changes take effect. This can be done via the Notes Admin client (**File > Tools > Server Administration > Console > etc.**).

11.2.0.0 INTELLIWATCH PAGING

11.2.1.0 Overview

IntelliWatch Paging (which is handled by the Message Center Gateway) adds a powerful notification feature to Monitor. As with SNMP messaging, it allows you to be notified of a Notes crash *with no Notes dependencies*.

Although the Message Center Gateway runs on both AIX and NT, IntelliWatch Paging is available only via a Gateway running on NT. However, as long as the Gateway is running on NT, pages can be sent from all platforms, either as a Command invoked by a Trigger, or as a command-line utility.

IntelliWatch Paging handles requests sent to the Message Center Gateway from the following Monitor components:

- Triggers and Commands
- Advanced ServerWatch
- Send Page Utility (see interface, below left)
- Command-line Paging Utility

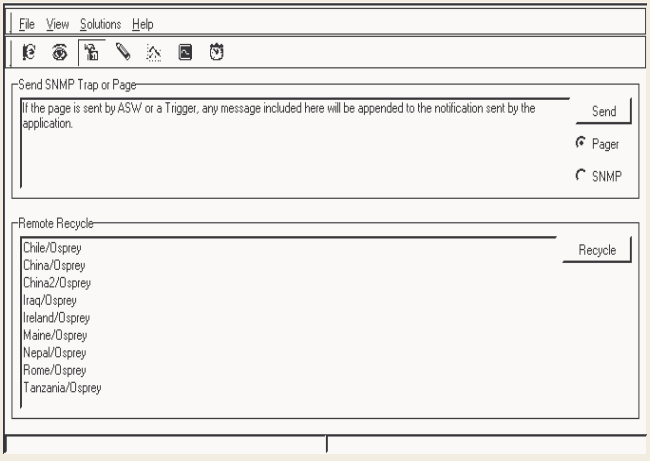

11.2.2.0 System Requirements

- Windows NT 4.0 workstation or server
- modem (connected to a landline)
- TCP/IP network connection

11.2.3.0 Modem Speed

Modem speed is not critical, since pager messages are relatively short and do not take long to transmit. *Far more important is modem quality and reliability.*

11.2.4.0 Supported Paging Protocols

- TAP (Telocator Alphanumeric Protocol)
- SMS (Short Messaging Service). This European protocol routes messages to mobile phones which are capable of receiving them.

Remember ...

Check with your paging service to see what protocols they support, and to obtain the dialup numbers for data-paging transmissions.

11.2.5.0 Does IntelliWatch Paging require Notes?

No. *However*, on a machine with Notes installed (whether a Domino server or just a Notes client), you can install the Pinnacle Console and use the Parameter Configuration Utility to access and modify the registry settings for the IntelliWatch Paging.

These settings can also be accessed via the machine's NT registry. The advantage of using the Parameter Configuration Utility is that you are accessing IntelliWatch settings only, and cannot change registry settings for another application by accident.

11.2.6.0 Before installing the Message Center Gateway

Before installing the Message Center Gateway, please have on hand the following information.

- pager number
  (phone number of paging company)

Remember ...

Your voice number won't work for sending pages. If a prefix must be dialed to get an outside line, be sure to enter that number—followed by a comma—before your pager number.

- pager IDs
  (comma-delimited, if more than one)
- confirmation from paging company that they have a modem pool, and can accept data messages

11.2.7.0 After installing the Message Center Gateway

- Verify that Paging has been selected where appropriate in your Notification Profiles in Advanced ServerWatch.
- Test Paging, using the command-line version of the Paging Utility (see *"iwpage" on page 371*).

11.2.8.0 Where to install the Message Center Gateway

IntelliWatch Paging is part of the Message Center Gateway, and, as such can be installed on an NT Workstation without Notes. *However*, the configuration of the machine where you install the Message Center Gateway influences how easily you can access *Paging Server* settings. Please note the advantages/disadv antages of the following locations:

11.2.8.1 System WITHOUT a Notes Client (Not Recommended)

- NT Workstation or Server,
- No Notes installation

All configuration must be done *on the Message Center Gateway itself*, by accessing the following registry key on NT:

HKEY_LOCAL_MACHINE\SOFTWARE\Candle\IntelliWatch\Pinnacle\Paging Server

On UNIX systems, see the corresponding entry in iwpinn.ini (located in the Domino data directory).

11.2.8.2 System WITH a Notes client

- NT Workstation or Server,
- Notes Workstation installed

- Pinnacle stand-alone client installed
  - Configure the settings (locally) using the Pinnacle Console.

11.2.8.3 On a machine running Domino Server (Recommended Setup)

- NT Workstation or Server,
- Domino Server installed
  - If the server is running, you may access the *Paging Server* settings remotely (or locally) using the Parameter Configuration Utility.

Configuration settings take effect after the IntelliWatch Messaging Center (NT Service) has been stopped and restarted. This can be done through the NT Service Control Panel, or by rebooting the machine.

11.2.9.0 IntelliWatch Paging: Configuration Settings

The following settings can be configured on all platforms via the Parameter Configuration Utility by going to **Pinnacle > Paging Server** (see *"Chapter 11, Parameter Configuration Utility"*):

- **Comm Port**
  - COM port for paging modem
- **Paging Number**
  - paging central telephone number
- **Pager ID**
  - pager ID
- **Init String**
  - initialization string for pager (default = empty string)
- **Termination String**
  - string that tells modem to hang up (usually ATH)
- **Dial Type**
  - tone or pulse
- **Password**
  - any password required to access paging service
- **Dial Attempts**
  - number of times to retry paging command, if unsuccessful
- **Port Speed**
  - baud rate of paging modem (default = 1200)
- **TCPIP Port Number**
  - TCP/IP port of NT system running Message Center Gateway
- **Max Page Length**
  - maximum character length of page (default = 160 characters). Must not exceed 240 characters with spaces (a limitation imposed by the TAP protocol).
- **Timer**
  - amount of time to wait before moving messages with a *new paging number* to top of queue

Once the Message Center Gateway connects to a paging number, as long as there are messages for that number, they take priority over messages to be sent to other paging numbers. Pages are sent to the other numbers only after the modem hangs up from the current number.

The Timer parameter forces the modem to hang up after a certain number of seconds. Pages are then sent to alternate numbers for the amount of time specified by the

Timer parameter, at which point the Message Center Gateway cycles through any additional numbers (or back to the numbers dialed earlier).

- Debug
  - If enabled, prints debugging information to **iwpage.log** in the Candle log directory.

    ENABLED=1; DISABLED=0

11.2.10.0 *Paging* vs *Paging Server* Settings

Earlier versions of IntelliWatch sent alphanumeric pages via a so-called Paging Server (that ran as an NT Service).

Although IntelliWatch encorporates paging functionality into the Message Center Gateway, and there is no longer a separate Paging Server, certain configuration settings still refer to *Paging* and the *Paging Server*.

This not only offers the advantage of familiarity, it also makes clear which Message Center Gateway settings relate to paging.

What's intended by *Paging* and *Paging Server* settings, respectively?

In short, *Paging Server* settings govern *the Message Center Gateway itself*. That is to say, these settings are the defaults used to send pages, unless they are overwritten by *Paging* settings on individual paging-client systems.

Paging settings should be thought of as local, that is they are relevant only for the system on which they reside.

These local settings override the settings on the Paging Server (NT system running the Message Center Gateway).

11.2.10.1 Parameter Configuration Utility

On all platforms, *Paging* and *Paging Server* settings can be configured using the Parameter Configuration Utility, accessed via the Pinnacle Console..

- *Paging* settings are accessed by navigating to**Monitor > Paging**
- *Paging Server* settings are accessed by navigating to Pinnacle **> Paging Server**

11.2.10.2 Configuring via the NT registry

On NT 4.0 and Windows 2000, these settings can be accessed via the registry, at: HKEY_LOCAL_MACHINE\SOFTWARE\Candle\IntelliWatch\Monitor\Paging and ...\Pinnacle\Paging Server, respectively.

11.2.10.3 Configuring via .ini files on UNIX

Paging settings are located in *iwmon.ini*
Paging Server settings are located in *iwpinn.ini*

11.2.10.4 Example of *Paging* settings: Changing the Message Center Gateway

Server_456 has served as the Message Center Gateway for pages sent from ServerGroup1. You now want those pages sent to Server_789.

TO CHANGE THE DESTINATION PAGING SERVER:

1. Click on the  toolbar icon; alternatively, go to **Solutions > Parameter Configuration** via the drop-down menus.
2. Using the combo box at the upper left of the interface, select one of the servers in ServerGroup1.
3. Click on the  icon (to the right of the combo box).

Current settings for installed components are retrieved and displayed in the left-hand pane.

4 In the left-hand pane, navigate to **Monitor > Paging > Paging Server**.

Click on the +'s to expand folders until you see the list of current settings.

5 In the left-hand pane, select the **Paging Server** parameter.

A value table is displayed in the right-hand pane

6 In the right-hand pane, select the parameter.

The background of the parameter changes from white to blue, and the edit icon is activated.

7 To change the parameter's value, click on the icon; alternatively, double-click the parameter in the right-hand pane of the Console.

An Edit Parameter dialog is displayed, with fields for both the name of the Paging Server and the list of servers for which the value should be changed. For an example of the dialog, see Figure 11-2 on page 351.

8 In the first field, enter Server_789 as the new Paging Server.

9 In the bottom field, enter ServerGroup1.

10 Click OK.

This changes the parameter's value in **iwparam.nsf** on all servers in that group.

Unless and until this parameter is changed for these servers, pages sent from ServerGroup1 will be routed to Server_789, instead of Server_456.

11.2.10.5 Example of *Paging Server* settings: Changing the COM port

Server_123 is the Paging Server for the ABC Company.

The modem used to send pages has always been attached to COM1 of Server_123, but in future will be attached to COM2 on that system.

In this case, settings must be changed on the Paging Server itself (under **Pinnacle > Paging Server > Comm Port**).

1 Click on the toolbar icon; alternatively, go to **Solutions > Parameter Configuration** via the drop-down menus.

2 Using the combo box at the upper left of the interface, select Server_123.

3 Click on the icon (to the right of the combo box).

Current settings for installed components are retrieved and displayed in the left-hand pane.

4 In the left-hand pane, navigate to **Pinnacle > Paging Server**.

Click on the +'s to expand folders until you see the list of current settings.

5 In the left-hand pane, select **Comm port**.

A value table is displayed in the right-hand pane

6 In the right-hand pane, select the parameter.

The background of the parameter changes from white to blue, and the edit icon is activated.

7 To change the parameter's value, click on the icon; alternatively, double-click the parameter in the right-hand pane of the Console.

An Edit Parameter dialog is displayed. In this case, the dialog has only one field, since the parameter being changed is effectively local. For an example of the dialog, see Figure 11-1 on page 351.

8 Enter COM2 in the dialog field.

9 Click OK.

This changes the parameter's value in **iwparam.nsf** on the Paging Server only.

10 Stop and restart the IntelliWatch Messaging Center service to read in the new settings.

Until the IntelliWatch Messaging Center service is stopped and restarted, pages continue to be sent to COM1!

Unless and until this parameter is changed on the Paging Server, pages will be sent via COM2 instead of COM1.

11.2.11.0 Sending a page

1 Click on the  toolbar icon at the Pinnacle Console.

The interface used to send pages and SNMP traps, and for recycling servers is displayed.

2 Select the Pager radio button (to the right of the text box).

3 Type in your message in the text box.

4 Click the Send button.

11.2.12.0 Debugging Paging

To access **iwpage.log**, go via the Start menu to:

Start > Programs > IntelliWatch > Monitor > Messaging Center Log

Q: When should I leave Debugging off?

If you're consistently receiving your pages, leave Debugging off (the default), as this limits the growth of **iwpage.log**.

With Debugging *disabled*, only essential information is logged:

- time the page was sent
- client (originating the page)
- telephone number
- pager ID
- message
- routing info (the pager ID)

Q: When should I turn on Debugging?

Should problems arise, turn Debugging on (set the value to 1). You can do this via the Parameter Configuration Utility (discussed above), or via the NT registry at:

HKEY_LOCAL_MACHINE\SOFTWARE\Candle\IntelliWatch\Pinnacle\Paging Server\Debug

Disabled=0; Enabled=1.

Enabling Debugging causes the following *additional* fields to be published:

- dial type (tone or pulse)
- initialization string
- comm port
- port speed
- modem retries
- session retries
- termination string
- "Transmitting" and "Received" information

With Debugging on—and if the page is successfully sent by the Message Center Gateway—**iwpage.log** information is published under the headings:

- Initializing Session: DEBUG
- Transmitting Page Request

With Debugging on—*and if the page is NOT successfully sent by the* Message Center Gateway—**iwpage.log** information is published under the headings:

- Initializing Session: DEBUG
- Could Not Fulfill Page Request, Check Settings
- To view a complete list of pager error messages, see *"Paging Error Messages" on page 391*.

11.2.13.0 Stopping and Restarting the Message Center Gateway

As noted above, changes made to the Message Center Gateway's *Paging Server* settings take effect only after the corresponding NT Service has been stopped and restarted. To do this without rebooting the machine:

1 Select **Start...Settings...Control Panel**.

2 Click on Services.
3 Select the IntelliWatch Messaging Center from the list.
4 Click on the Stop button.
5 Make any desired configuration changes.
6 Click on the Start button.

11.3.0.0 REMOTE RECYCLE UTILITY

11.3.1.0 Overview

Occasionally, Notes administrators need to recycle a server remotely. The Remote Recycle Utility was designed for just that purpose.

In fact, the Remote Recycle Utility allows you to simultaneously recycle several servers.

11.3.2.0 Requirements

For Remote Recycle to work, both of the following must be true:

- You currently have TCP/IP connectivity to the servers to be recycled.
- IntelliWatch Monitor is installed on the selected servers.

11.3.3.0 Procedure

1 Click on the  toolbar icon at the Pinnacle Console.

The interface used to send pages and SNMP traps, and for recycling servers is displayed.

2 Select the server(s) you want to recycle from the list.
3 Click the Recycle button.

11.3.4.0 Steps in process

Remote Recycle takes the following *progressive* steps:

- **STEP 1:** attempts to shut down the Domino server and all add-in tasks, then restart the server.
- **STEP 2 (if STEP 1 fails):** attempts to kill the Domino server process and restart it.
- **STEP 3 (if STEP 2 fails):** reboots the operating system and brings up the Notes server.

11.4.0.0 REPLICATION CHECK

11.4.1.0 Overview

Replication Check searches for replica databases on two servers, then checks the databases selected for New Documents, Differences and Errors.

Use Replication Check to verify and troubleshoot replication.

11.4.2.0 Requirements and preconditions

Replication Check is available only at the stand-alone Pinnacle client. The Solution does not appear when the Pinnacle Console (**console.nsf**) is accessed via a browser.

The Replication Check user interface is supported on NT 4.0 and Windows 2000, but not on Windows 95/98. (See also Table 1-2 on page 24.)

- For each Source Server in the list, there must be:
  - an entry for the Source Server in the Servers view of the NAB
  - a connection document in the Connections view of the NAB with either:

  - the Source Server as *source*, and replication set up such that changes can flow from the Source Server to the Destination Server (that is, the replication direction is one of Push-only, Push-Pull, or Pull-Pull), or
  - the Source Server as a *destination*, and replication set up such that changes can also flow from the Destination Server to the Source server (that is, the replication direction is one of Pull-only, Push-Pull, or Pull-Pull).

- Similarly, for each Destination Server, there must be:
  - an entry for the Destination Server in the Servers view of the NAB
  - a connection document in the Connections view of the NAB with either:
    - the Source Server as the *source*, the Destination Server as destination, and replication set up such that changes can flow from Source to Destination (that is, the replication direction is one of Push-only, Push-Pull, or Pull-Pull), or
    - the Destination Server as *source*, the Source Server as destination, and replication set up such that changes can also flow from Source to Destination (that is, the replication direction is one of Pull-only, Push-Pull, or Pull-Pull).

For all connections, the following must be true:

- Replication is one of the tasks to be performed on that connection.
- The replication schedule for the above connection is enabled.

With bi-directional replication, the set of notes compared per direction can be different due to different settings for each of the replicas.

11.4.3.0 Console

Selecting the  toolbar icon brings up the Replication Check user interface.

11.4.3.1 Combo boxes

- **Select Source Server**
- **Select Destination Server**

The list of available servers is retrieved from the local NAB, and is the same in both combo boxes. Databases are recovered only if connection documents are present.

11.4.3.2 Menu items

- **View > Show Note ID**

  ID in the context of the database.

- **View > Show Unique ID menu option**

  ID consistent across all replicas.

11.4.3.3 Check boxes

- **Retrieve Extended Information**

  The Note ID, followed by the document title (New and Different documents).

  Deselect to return the Note ID only (which may speed up processing).

- **Use Replication Formula**

  If selected, the subset of documents returned is governed by applicable replication formulae.

11.4.3.4 Icons

-  **icon initiates replica search**

  (See *Figure 11-3* to view search results.)

-  **icon compares found replicas**

  (See *Figure 11-4* to view the results of a comparison.)

FIGURE 11-3: Results of replica search

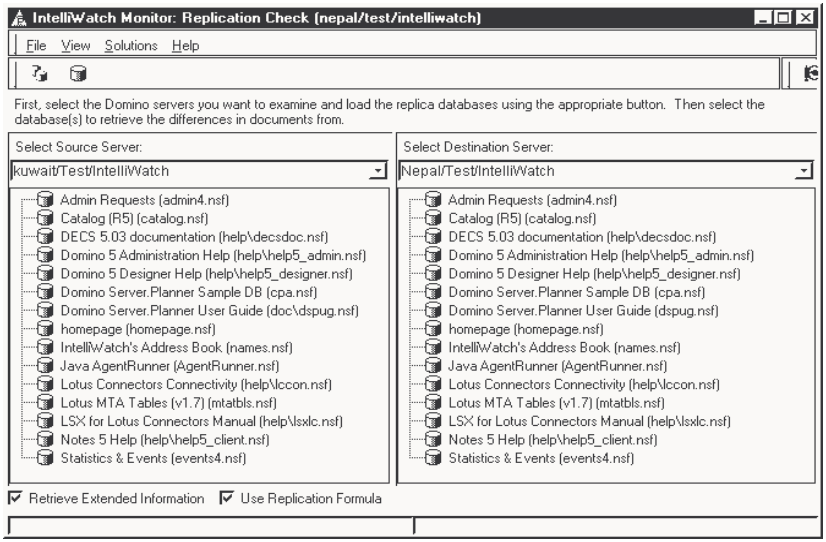

FIGURE 11-4: Results of database comparison

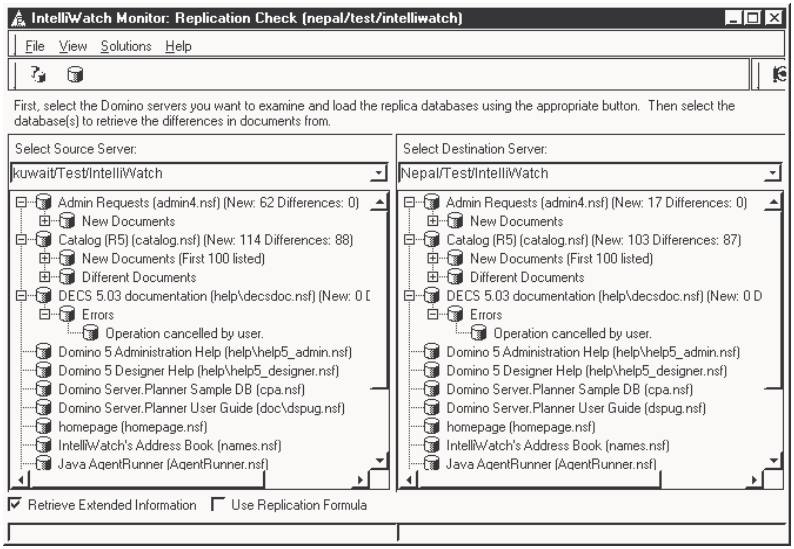

11.4.4.0 What documents are returned?

Replication Check does not return documents it deems identical. (See *"Replication settings" on page 362* for the basis of that determination.) Documents are returned if—and only if—they fall into one of the following categories (reported individually by database):

- **New documents**

  New Documents are found in the replica on one server or the other—*but not both*.

  A New Document in one replica database has a Unique ID (UNID) that is not shared by *any* documents in the other replica examined by Replication Check.

  Administrators can use the list of New Documents returned by Replication Check to fine-tune replication settings for the databases in question.

- **Different documents**

  Different Documents are found in *both* replica databases.

  A Different Document in one replica has the same UNID as the corresponding document in the other replica.

  Their having the same UNID means that, at one time, the two documents were identical. Over time, however, changes were made to one or both documents.

  Only the one modified most recently is displayed.

- **Errors**

  May show up in only one window (with the opposite window displaying a corresponding Error stub), or in both the Source and the Destination windows, depending on the particular error condition detected.

  Errors are also reported if the current operation is cancelled. While no actual error may have been found in a document, partial results are not returned as they may be inaccurate.

11.4.5.0 Understanding the display

11.4.5.1 Color-coding

Replication Check uses color-coding to make clear the three possible outcomes of the comparison.

- **Replicas in sync**

  If no Differences and no Errors were found by Replication Check, replicas are deemed to be *in sync*.

  No documents are displayed, and the color of the corresponding entry changes to green.

The set of documents compared depends on the Notes replication settings. Databases are considered 'in sync' if the documents compared are deemed identical. Documents which are not compared may NOT be in sync.

- **Differences between replicas**

  The most recent version is displayed (reported individually by database).

  The color of the corresponding entry changes to yellow, and becomes a folder with descendant nodes displaying the different document(s).

- **Errors**

The color of the corresponding entry changes to red, and becomes a folder with descendant nodes displaying the error message(s).

11.4.5.2 Note ID vs UNID

The drop-down menus allow you to display returned documents based on either the Note ID or the UNID. What is the significance of these two identifiers?

A document's Note ID is a unique identifier, but *only* within the context of *that copy* of the database. Note IDs are *not* consistent across replica copies of a database.

This is true even immediately after a document has been replicated from, say, ReplicaDbA to ReplicaDbB (and before any changes have been made to *either* copy of the document).

Although the two copies in this example have different *Note IDs*, their UNID *is* identical. And that remains true, even *after* changes have been made to one or both copies of the document.

You may want to display documents in Replication Check by Note ID, however, if you work with returned documents in a view that includes that identifier.

To switch to displaying documents by UNID, simply go to **View > Show Unique ID** via the drop-down menus.

11.4.6.0 How document identity is determined

Since replica copies of a document maintain their UNID, how does Replication Check determine document identity (or establish precedence between copies of the document exhibiting differences)?

Document identity is determined by a third identifier, called the Originator ID (or OID, also known as the Document Version ID).

11.4.6.1 What makes up the OID?

The OID includes the UNID, but adds a Sequence Number and a Sequence Time/Date. Replica documents are considered to be different if their Sequence Number and/or Sequence Time/Date are different.

For information on the details of Replication, including the various ID types, please see your Notes documentation.

11.4.7.0 Replication settings

Replication Check can take into account some—but not all—Notes replication settings:

The following menu references assume you have already accessed your replication settings via a Notes client. (... > stands for **File > Replication > Settings**.)

- **Formulas and document categories**

  go to **... > Advanced** (button) **> Replicate** (group) **> etc.**

- **Purge interval**

  go to **... > Space Savers** (button) **> Remove documents not modified in the last:** (field).

- **Cutoff time**

  go to **... > Other** (button) **> Only replicate incoming documents saved or modified after:** (field).

Replication Check does not take into account settings configured at **... > Advanced** (button) **> Replicate a subset of documents:** (area).

11.4.8.0 Performing a Replication Check

TO PERFORM A REPLICATION CHECK:

1 Select the toolbar icon to bring up the Replication Check user interface.

2 Select the Source server from the drop-down list on the left-hand side of the Replication Check window.

3 Select the Destination Server from the drop-down list on the right.

4 Click the . icon.

A list of databases with the same replica ID on the two servers is created. While retrieval is taking place, a progress bar is displayed.

FIGURE 11-5: RepCheck progress bar

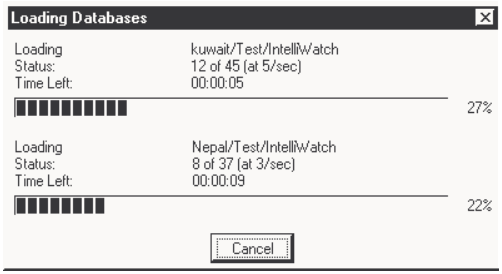

5 From the generated list (for an example, see *Figure 11-3*), select the database(s) whose replicas you want to check for differences, errors, and so on.

6 Click the toolbar icon to compare databases.

Selecting the database in either pane of the user interface causes its counterpart on the other server to be selected in the opposite pane.

7 Examine the created list for New Documents, Differences and Errors.

See *Figure 11-4* for the results of a sample comparison.

11.5.0.0 SEND SNMP TRAP UTILITY

11.5.1.0 Overview

SNMP (Simple Network Management Protocol) allows you to communicate with network devices, to gather information, or to change settings for that device.

The IntelliWatch Send Trap Utility allows you to send traps from your Admin workstation to your SNMP Network Management Station.

11.5.2.0 Interface

The user interface is the same one used by both Paging and Remote Recycle.

11.5.3.0 Prerequisites

For the IntelliWatch Send Trap Utility to function properly, SNMP must be properly configured both on your workstation and at the Network Manager.

For assistance, please contact your local SNMP support team, or the software manufacturer of the Network Manager in use at your company.

11.5.4.0 Procedure

1 Click on the toolbar icon at the Pinnacle Console.

The interface used to send pages and SNMP traps, and for recycling servers is displayed.

2 Select the SNMP radio button (to the right of the text box).

3 Type in your message in the text box.

4 Click the Send button.

11.5.4.1 Trap priority

The Priority field is optional. If left blank, the value is taken from the registry setting on the Admin workstation: HKEY_LOCAL_MACHINE/SOFTWARE/

```
Candle/IntelliWatch/Monitor/SNMP/Defaul
  t Priority)
```

Since the Pinnacle Console does not run on UNIX systems, there is no corresponding setting on Solaris and AIX.

Chapter

Command-line Utilities

12

Pinnacle Enterprise includes several command-line utilities that can be launched from a command prompt, or in a batch file.

To obtain usage information for a utility, type the name of the executable at a command prompt:

> iwconcmd

All IntelliWatch command-line utilities make use of the same conventions:

<>designate a required parameter

[] designate an optional parameter

Quotation marks (“”) must be used for *any argument containing spaces*. For example, if the parameter in question is the name of a server,

NAME_OF_SERVER need not be enclosed in quotes; NAME OF SERVER requires them.

Remember ...

*Commands entered at the server console must be prefaced with “l[oad]”. The examples in this chapter assume utilities are being launched at a command prompt (as they **must** be on UNIX).*

12.1.0 IWALERT

Sends messages to ASW Hub Server(s), for viewing at the Pinnacle (or ASW) Console.

iwalert <"/M:*Message*"> [/H:*HubServer*] [/P:*Priority*]

12.1.1 Variable definitions:

Variable	Definition
message	message to be sent
HubServer	name of Advanced ServerWatch Hub
priority	message priority

12.1.2 Example:

To send the priority 2 message "Srvr Not Responding" to HubServer, enter:

```
iwalert "/M:Srvr Not Responding" /P:2 /H:HubServer
<CR>
```

12.2.0 IWCONCMD

Used to execute a Notes server console command in a batch file or from a command line:

iwconcmd [*command*] [*server*]

12.2.1 Variable definitions:

Variable	Definition
command	command you want to be executed
server	name of server on which to execute the command

12.2.2 Examples:

To load the Router task on NOTES_SERVER, enter:

```
iwconcmd "load router" NOTESSERVER
<CR>
```

To show Domino server tasks currently running on NOTESSRV1, enter:

```
iwconcmd "show tasks" NOTESSRV1
<CR>
```

12.3.0 IWEVENT

Used to send Events (notifications) to IW Message Center Gateway.

Events can be of three types:

- SNMP
- NTLog (NT only)
- Tivoli TEC Event

iwevent <[EventType]> </m:Message> [type-specific parameters]

12.3.1 Variable definitions:

Variable	Definition
EventType	one of three types: SNMP, NTLog, TEC
Message	string to be sent with the Event
type-specific parameters	type-dependent parameters may or may not be required (type "iwevent [EventType]" for help)

12.3.2 Example:

To send an SNMP trap to your Network Manager from ServerABC, with a priority of 6 and including the message string "ServerABC needs SP update", enter:

```
iwevent snmp "/m:ServerABC needs SP update" /f:serverabc /p:6
<CR>
```

In this example, the /f: and /p: parameters are type-specific. Type "iwevent SNMP" to obtain usage for the SNMP event.

The same events listed above can also be sent via the Message Center Gateway by using the following type-specific utilities:

- **iwsnmpevt** (see also ***12.9.0***, below)
- **iwntlogevt** (see also ***12.5.0***, below)
- **iwtecevt** (see also ***12.11.0***, below)

Use these latter utilities as you would **iwevent**; just eliminate the Event Type.

12.4.0 IWMAIL

iwmail <"*Message*"> [/F: *From*] [/T: *To*]
[/P: *Password*] [/CC: *To*]
[/S: *Subject*] [/A: *Attached File Name*]
[/I: I*mported File Name*]

12.4.1 Example:

To send Frank Smith the message "Remember the meeting at 2pm", Mark Jones would enter:

```
iwmail "/M:Remember the meeting at 2pm" /F:Mark_Jones /T:Frank_Smith
<CR>
```

12.5.0 IWNTLOG

iwntlog <"*Message*"> [-t *Type*] [-c *Category*] [-n *Computer Name*]

12.5.1 Example:

To send the NT Event Log the message "Potential database corruption on Svr123", enter:

```
iwntlog "Potential database corruption on Svr123"
<CR>
```

Usage: IWNTLog "Message" [-t Type] [-c Category] [-n Computer Name]

Example: IWNTLog "Access denied" -tF -cA

Options:

-t Type is one of I (Information)

W (Warning)

E (Error)

S (audit Success)

F (audit Failure)

-c Category is one of G (General)

U (commUnication)

C (Corruption)

D (Database)

B (Database Corruption)

O (dOcument Corruption)

X (full Text Corruption)

M (Mail)

R (Replication)

E (rEsource)

S (Server)

A (server Access)

T (Template)

-n Computer Name is the name of a computer in UNC format.

12.6.0 IWPAGE

Current settings for the parameters in parentheses are contained in the registry of the paging server. These arguments should be included on the command line only if you want to override them. For that reason, they are listed below in parentheses rather than between '<>'s or '[]'s.

iwpage (/C:*computer*) (/W:*timeout*) (/L:*TCPIP_port*) (/R:*Retries*) (/P:*port*) (/B:*baud_rate*) (/N:*number*) (I:*pager_id*) (/S:*init_string*) (/T:*dial_type*) <"/M:*message*">

12.6.1 Variable definitions:

Variable	Definition
computer	paging server computer name (default = local host)
timeout	timeout in seconds for connecting to paging server
TCPIP_port	TCP/IP port of paging server
retries	how often to retry paging command if unsuccessful
port	COM port of paging modem
baud_rate	baud rate of paging modem (default =1200)
number	paging central number
pager_id	pager ID
init_string	initialization string for pager (default = empty string)
dial_type	specifies tone or pulse dialing
message	message to send to pager, enclosed in quotes

12.6.2 Example:

To send a "Test page" message to the Paging Server (pserver), enter:

```
iwpage /C:pserver "/M:Test Page"
<CR>
```

12.7.0 IWRECYCL

The standard method for obtaining usage help for an IntelliWatch command-line utility is to type in the name of the utility without arguments.

Do not do this with the iwrecycl Command!

At an Admin workstation, typing in iwrecycl results in a "command not recognized" message; on a server, typing in iwrecycl launches the utility and recycles the local server!

iwrecycl [/T:*frequency*] [/R:*retries*]
[/W:*stop_time*] [/P:*start_time*]
[/M: 1 0r 0] [/S:*server_name*]
[/D:*data_directory*] [/B:*batchfile*] [/F:*batchfile_timeout*]
[/L:*trap_priority*]
[/V:*notification_level*] [/A:Notes_API_tasks] [/Z: 1 or 0]

12.7.1 Variable definitions:

Variable	Definition
frequency	interval at which to check for server response
retries	number of times to check before deciding server is down
stop_time	how long to wait for Notes server to stop Same as Terminate Notes Time in Recycle section of NT registry/iwmon.ini.
start_time	how long to wait for Notes server to start Same as Wait for Server to Restart in Recycle section of NT registry/iwmon.ini.
/M	enables/disables paging 1=paging enabled, which sends a page before rebooting. 0=paging disabled.
server_name	name of server
data_directory	path to Domino data directory
batchfile	batchfile to run
batchfile_timeout	maximum time, in seconds, to wait for batch file to complete
notification_level	verbose, limited, or none 'Limited' notifies you whether the server was successfully restarted or not.
Notes_API_tasks	list of add-in tasks to recycle with server Same as the Notes Programs listing in Recycle section of NT registry/iwmon.ini.
/Z	indicates machine should/should not be rebooted if Notes server cannot be restarted 1=reboot, 0=do not reboot.

Notification Level options are set in the Recycle section of:

– the registry (NT) at

HKEY_LOCAL_MACHINE/SOFTWARE/Candle/IntelliWatch/Monitor/Recycle/Notification Level

– *iwmon.ini* (UNIX) at

```
[Recycle]
Notification Level
```

- **None** means Recycle sends no notifications.
- **Limited** means Recycle uses the enabled forms of notification to inform you whether or not the server was successfully restarted.
- **Verbose** means Recycle uses the enabled forms of notification to inform you at each significant stage of the Recycle process.

"Enabled forms of notification" refers to the options listed in the Recycle section referred to above.

Do not use the Verbose setting unless instructed to do so by IntelliWatch Customer Support, since it produces a significant number of notifications.

12.7.2 Example:

The following command line loads Recycle on MYSRV, checks for server response every five minutes (300 seconds), and launches the backup.bat file:

```
iwrecycl /s:MYSRV /b:backup.bat /f:300
<CR>
```

12.8.0 IWRUNAGT

iwrunagt </D:*database*> </A:*agentName*>

12.8.1 Variable definitions:

Variable	Definition
database	filename of database containing agent
agentName	name of agent to run

12.8.2 Example:

To run the "Release Dead Mail" agent on mail.box, enter:

```
iwrunagt /D:mail.box "/A:Release Dead Messages"
<CR>
```

12.9.0 IWSDTRAP

iwsdtrap <"*trap_string*"> [/P:*trap_number*] [/E:*enterprise_id*] </S:*server*>

12.9.1 Variable definitions:

Variable	Definition
trap_string	message to be sent with trap, enclosed in quotes
trap_number	trap priority (default = 4)
enterprise_id	enterprise ID of SNMP trap
server	name of server sending trap

12.9.2 Example:

To send a "Srvr Not Responding" message from MYSRV to the SNMP Network Manager, enter:

```
iwsdtrap "Srvr Not Responding" /S:MYSRV
<CR>
```

12.9.3 IWSNMPEVT

In addition to **iwsdtrap**, IntelliWatch provides another utility for sending SNMP traps, namely **iwsnmpevt**. The difference between these two utilities is that **iwsdtrap** sends traps via local SNMP functionality, whereas **iwsnmpevt** sends the trap to the system running the IntelliWatch Message Center Gateway.

Parameter usage is the same for both utilities, with one exception:

- The “trap_string” must be preceded by /M: when using **iwsnmpevt**.

12.10.0 IWSLEEP

iwsleep <*time*>

12.10.1 Variable definitions:

Variable	Definition
time	number of seconds to wait before executing next action

12.10.2 Example:

For a 20-second pause, insert the following line in the batch file:

```
iwsleep 20
```

12.11.0 IWTECEVT

iwtecevt </C:eventClass> </M:Message> [/H:hostname] [/S:severity] [/T:status] [/U:subSource] [/N:<slotname>=<value>;...<slotname>=<value>;]

12.11.1 Variable definitions:

Variable	Definition
eventClass	
Message	Message to be sent should be in quotes if it contains spaces.
hostname	IW Message Center hostname (example: MyServer.MyCompany.com).
severity	
status	
subSource	
slotname	

12.11.2 Example:

To send a generic test message to Tivoli:

```
iwtecevt /C:IW_GENERIC_EVENT "/M:This is a test"
```

Chapter

13 In the Latest Maintenance Release

This chapter covers features and fixes included in the latest maintenance release of IntelliWatch, 6.00.27.40.

Depending on the version of IntelliWatch running in your environment, the information in this chapter may supersede product information in other sections of the manual.

Chapter Contents

13.1.0.0 FUNCTIONALITY UPDATE

13.1.1.0 Blackberry servers

Though no product changes were involved, we're pleased to pass on the news that some customers have—starting with Pinnacle 99!—been using IntelliWatch Monitor Triggers to monitor Blackberry servers.

13.1.2.0 Solaris on Domino 6 and above

Due to a non-IntelliWatch issue, the product could not be installed on Solaris systems running Domino 6 and higher. In cooperation with IBM, a workaround for the issue was developed. For details, see *"On Solaris systems running Domino 6 and above" on page 35*.

13.1.3.0 Product fixes

For a list of product fixes in this maintenance release (by component), see the update.txt file which can be found on the installation media.

This file also may be found at the default location of update.txt on Windows systems, and in the directory where IntelliWatch installation files are unpacked on UNIX.

Appendix

A

Definitions of PM Statistics Variables

The following table contains definitions for the variables used with PM statistics.

You can obtain similar information by placing the mouse cursor over the field in question.

Table A-1. Statistic Variable Definitions (Sheet 1 of 2)

Statistic Type	Element	Definition
Add-in/View Performance	Add-in task	Notes Add-in task for which you want to calculate a performance statistic.
Add-in/View Performance	Database	Notes database for which you want to calculate a performance statistic.
Add-in/View Performance	Performance Type	Enter Performance Type you want to count to enable the correct fields below.
Add-in/View Performance	User	Notes User for whose actions you want to generate performance statistics.
Add-in/View Performance	View	Database view for which you want to calculate a performance statistic.
Average	Add	Add statistics not in the list.
Average	Select Statistics	Select statistics to be averaged.
Delta	Add	Add statistics not in the list.
Delta	Select Statistic	Statistic to be reset.
Difference	Add	Add statistics not in the list.
Difference	Select First Statistic	Statistic from whose value you want to subtract.
Difference	Select Second Statistic	Statistic whose value is to be subtracted from First Statistic.
Mail Delivery Threshold	Threshold	Occurrence threshold (count instances in excess of this number).
Mail Size	Maximum/Average	Statistic based on maximum and average size of all mail routed/delivered.
Mail Size	Report Threshold	Minimum size reported (not valid for Maximum/Average).
Mail Size	Server Level Report	Statistic based on maximum and average size of mail routed/delivered. (Broken down by destination server (local server excepted).
Mail Size	User	User whose mail attachment size you want to count. Leave blank for all users

Table A-1. Statistic Variable Definitions (Sheet 2 of 2)

Statistic Type	Element	Definition
Mail Size	User Level Report	Statistic based on maximum and average size of mail routed/delivered. (Broken down by user.)
NT Performance Counter	Counter	Object component to be counted.
NT Performance Counter	Instance	Specific instance of object to be counted.
NT Performance Counter	Object	NT Counter object for which to generate a statistic.
NT Performance Counter	Unlisted	Add instances not in list.
Replication Delay	Replication Database	Database for which to count replication time.
Server Event Count	Add-in Task	Notes Server Add-in task whose actions are to be counted.
Server Event Count	Database	Database against which counted actions are taken.
Server Event Count	Internal Server Events	Event you want to count.
Server Event Count	User	Notes User whose actions you want to count.
Summation	Add	Add statistics not in list.
Summation	Select Statistics	Select statistics to be summed.

Appendix

IntelliWatch Keywords

B

Usage:

IntelliWatch keywords can be passed **1)** as arguments to a Start Program Command (to set a variable); **2)** in the message field of Triggers.

In either case, enclose the keyword in angle brackets <>.

Case-sensitivity: UNIX vs NT

On NT, keywords are not case-sensitive; on UNIX, however, they are. On UNIX systems, therefore, DATABASE and database are not the same.

The same is not true for flags. On both Windows and UNIX systems, /m: is the same as /M:, for example.

Compound Triggers and Keywords

For compound Triggers you must append 1 or 2 to the keyword depending on the associated condition. (See *"Compound Trigger" on page 83*.)

Table B-1. IntelliWatch keywords (in alphabetical order) (Sheet 1 of 5)

Keyword	Trigger Type	Description
APPLICATION	Application	Name of the application being monitored.
ACTIVITY_CONDITION	Database Activity	>,<, Range, increases by.
ACTIVITY_SCOPE	Database Activity	Activity for last day, week, month.
ACTIVITY_TYPE	Database Activity	Uses, Reads, or Writes.
ACTIVITY_VALUE	Database Activity	Actual number of times database was used, read, or written to.
ACL_CONDITION	ACL History	Search for any changes made by listed users or groups, or changes made by users or groups not in the list.
ACL_USER_LIST	ACL History	List of users, servers, and groups to watch for in the ACL history.
ACL_VALUE	ACL History	The actual events added to the ACL history which match the condition.
APP_INSTANCES	Application	Number of instances of the application which are running.
APPLICATION_STATE	Application	< or > the specified number of instances.
AVAIL_ERROR_MSG	Availability	If the port or database did not respond, this variable contains an error message explaining it in more detail.
AVAIL_TYPE	Availability	Notes DB or TCP application.
DATA_DIR	All	The Domino data directory.
DATABASE	Database, Database Activity, ACL History, Availability, Document Count, Document Timeout, White Space, Replication Integrity, Replication Readiness	Name of the database named in the Trigger.
DATABASE_CONDITION	Database	Contains, Does not contain, <, >, <>, = .
DATABASE_FIELD	Database	Name of the field in the database.

Table B-1. IntelliWatch keywords (in alphabetical order) (Sheet 2 of 5)

Keyword	Trigger Type	Description
DATABASE_OCCURRENCES	Database	Number of occurrences on the Condition tab of the Trigger.
DATABASE_OCCURRENCES_FOUND	Database	Number of occurrences encountered during a given monitoring cycle.
DATABASE_SCOPE	Database	Replaced with NEW, ALL, or BETWEEN (which requires specified begin and end times).
DATABASE_VALUE	Database	Actual event from the database that matched the Trigger.
DEST_ADDITIONAL_COUNT	Replication Integrity	Number of documents which exist in the replica server's database but not in the source server's database.
DEST_ADDITIONAL_IDS	Replication Integrity	IDs of documents which exist in the replica server's database but not in the source server's database.
DEST_MODIFIED_COUNT	Replication Integrity	Number of documents which exist in the replica server's database and exist in the source server's database, but which are not identical (some fields differ).
DEST_MODIFIED_IDS	Replication Integrity	IDs of documents which exist in the replica server's database and exist in the source server's database, but which are not identical (some fields differ).
DOC_COUNT_CONDITION	Document Count	<, >, <>, =, Range.
DOC_COUNT_VALUE	Document Count	Number of documents in database.
DOC_TIMEOUT_COUNT	Document Timeout	The actual number of documents that are older than the threshold.
DOCUMENT_ID	Document Timeout	Returns IDs of documents exceeding the specified timeout.
DOC_TIMEOUT_CONDITION	Document Timeout	Contains, Does not contain, <, >,<>, =.
ESCALATION_INDEX	All	Returns current Trigger Escalation stage.

Table B-1. IntelliWatch keywords (in alphabetical order) (Sheet 3 of 5)

Keyword	Trigger Type	Description
EXPLANATION	All	Causes the explanation field from the Trigger to be inserted into the message.
FILE	File	Name of the file in the file Trigger.
FILE_CONDITION	File	=, <>, <, >, Range.
FILE_TRIGGER_TYPE	File	Condition type of the File Trigger: Date/Time, Size, Exists, or Does not exist.
FILE_VALUE	File	Date/Timestamp or Size in bytes. (Applies to Date/Time and Size condition types.)
MAILUSER	Document Count	File user of the specified mail database.
MAIL_PROBE_REPLY_TIME	Availability	Time it took for an echo or response to come back from the address on the Internet.
MAIL_PROBE_SEND_TIME	Availability	Time it took mail probe to get to an SMTP server.
MESSAGE	All	Causes message specified in message field to be displayed.
NOTES_DB1	Database	All information matching the Trigger event field from the database. Useful in detecting which database is corrupt.
NOTES_DB2	Database	All information matching the Trigger event field from the second database. Useful in detecting which database is corrupt.
PROTOCOL	Database	Set to TCP/IP, NetBIOS, etc.
PORT	Database, Availability	The TCP application to monitor.
RANGE_FROM	Database	Range from value, if used.
RANGE_TO	Database	Range to value, if used.
REP_THRESHOLD	Replication Integrity	The value for the Replica Threshold field.

Table B-1. IntelliWatch keywords (in alphabetical order) (Sheet 4 of 5)

Keyword	Trigger Type	Description
READINESS_CONDITIONS	Replication Readiness	List of all errors that prevent replication between a source and replica server.
REPLICA_SERVER	Replication Readiness, Replication Integrity	Name of server on which replica database resides.
SERVER	Database	Name of the server.
SOURCE_ADDITIONAL_COUNT	Replication Integrity	Number of documents which exist in the source server's database, but not in the replica server's database.
SOURCE_ADDITIONAL_IDS	Replication Integrity	IDs of documents which exist in the source server's database, but not in the replica server's database.
SOURCE_MODIFIED_COUNT	Replication Integrity	Number of documents which exist in the source server's database and exist in the replica server's database but are not identical.
SOURCE_MODIFIED_IDS	Replication Integrity	IDs of documents which exist in the source server's database and in the replica server's database, but are not identical.
STATISTIC	Statistic	Name of the statistic.
STATISTIC_CONDITION	Statistic	=, <>, <, >, Range, Contains, Does not contain.
STATISTIC_VALUE	Statistic	Value of the statistic.
THIS_SERVER	All	The Domino server on which IntelliWatch Monitor is running
THRESHOLD	Application, Database, File, Statistic, White Space, Document Count, Document Timeout, Database Activity, Availability	The threshold value of the Trigger.
TIMEOUT	Availability	Actual time the database or port took to respond. (Zero does not always indicate success. Check avail_error_mesg to find out.)
USER	Database	Name of the user.

Table B-1. IntelliWatch keywords (in alphabetical order) (Sheet 5 of 5)

Keyword	Trigger Type	Description
USER_EXIT_STATUS	User	Return code from a user application.
VALUE	Database, File, Statistic, White Space, Document Count, Database Activity, Availability	The actual value that was found (can be used in place of file_value, statistic_value, etc.)
WHITE_SPACE_CONDITION	White Space	=, <>, >, <, Range.
WHITE_SPACE_VALUE	White Space	Amount of white space (for example., 20%).

Appendix

C

Paging Error Messages

The following table lists IntelliWatch Paging messages by error code.

Table C-1. Paging error messages (Sheet 1 of 2)

Code	Explanation
12	Unable to open COM port.
14	Unable to set baud rate on COM port.
16	Unable to set ASCII protocol parameters on COM port.
24	DosDevIOCtl Return code.
26	Unable to check receive queue.
36	Unable to write to COM port.
44	Failure detected by wait function.
54	Failure while closing COM port.
99	Modem not responding. Check modem connections. This message is presented after all attempts (2) have failed. Check that the modem is turned on, and properly connected to the configured PC COM port. Verify that serial port speed has been properly configured.
100	Unable to CONNECT after dialing. The local modem dialed out to paging central, but was unable to establish a connection with their modem. This failure occurs after three attempts to connect to paging central.
108	Invalid login prompt received from paging central.
116	No login prompt received from paging central.
124	Forced Disconnect received on login. For some reason, paging central has rejected or not replied to your login. Generally reflects a problem at paging central.
132	No reply received to login request. Generally reflects a problem at paging central.
140	Message "GoAhead" not found in response. If this happens more than once, speak with someone at your paging service; it is their responsibility to respond within a documented time interval.
164	Transaction abandoned by paging central. You likely supplied an invalid target pager ID.
172	No response received to message transmission. When this error occurs, IntelliWatch Monitor cannot determine whether or not the message was successfully delivered to the target pager. Something may have happened to the connection between the modems during or after transmission.

Table C-1. Paging error messages (Sheet 2 of 2)

Code	Explanation
180	Previously ACK'd data were later rejected.
188	No response received after Disconnect. It is possible (but not likely) that the page was not delivered.
196	Transaction aborted by paging central. Call your paging service provider.
200	Invalid number of parameters on invocation.
211	Invalid communications port name passed.
212	Message text too short.
223	Too many digits in dial number for paging central.
224	Too many digits in dial number for pager ID.

Appendix

IntelliWatch Databases: Template Usage and Replication

D

The following overview of template usage and replication is provided to assist you in the efficient management of IntelliWatch databases.

D.1.0 From media or via a template?

When Setup is run on the first server in a new environment, all IntelliWatch databases have to be created. Databases are either:

- directly copied from the media
  - all MAs
  - Command database
  - PM statistic definitions
  - Pinnacle Wizard
- or created via a template on the media
  - all other IntelliWatch databases

Two caveats must be mentioned:

- These guidelines apply with certainty only to the first installation in an environment.
- Whether they apply to subsequent installations depends on factors discussed at *D.1.1*.

Table D-1. IntelliWatch Databases: template usage and replication

Filename	Description	From Media/ Template?	Template Name(s)	Replicates?
analyzer.nsf	Analyzer Configuration	template	analyzer.ntf	no
console.nsf	Pinnacle Console	template	console.ntf	yes
iwadvsrv.nsf	Advanced Server MA	media		yes
iwasw.nsf	Advanced ServerWatch Configuration	template	iwasw.ntf	no
iwaswbk.nsf	Advanced ServerWatch Archiving	template	iwasw.ntf	no
iwcmd.nsf	Command Database	media		yes
iwcorrpt.nsf	Database Corruption MA	media		yes
iwmail.nsf	Temp Mail	template	mail45.ntf, mail46.ntf, mail50.ntf	no
iwmscc.nsf	Mail, Scheduling and CalConn MA	media		yes
iwnet.nsf	Internet MA	media		yes
iwparam.nsf	Parameters	template	iwparam.ntf	yes
iwpmstat.nsf	Statistic Definitions	media		yes
iwreplca.nsf	Replication MA	media		yes
iwreport.nsf	Analyzer Report Repository	template	iwreport.ntf	no
iwserver.nsf	Core Server MA	media		yes
iwsmtp.nsf	SMTP/MTA MA	media		yes
iwstats.nsf	Pinnacle Statistics Repository	template	iwstats.ntf	yes
iwstatus.nsf	Alternative Repository for some statistics	template	iwstatus.ntf	yes
pinWizrd.nsf	Wizard	media		yes
	Custom MAs	template	iwmon.ntf	yes

D.1.1 Environment-specifics

One factor more than any other influences whether the guidelines at *D.1.0* apply to your environment: the list of Pinnacle components installed on your Primary Server.

Consider the following example.

Example 1:

Assume an environment where all Pinnacle components have already been installed on ***YourPrimarySvr/YourCompany***—with the following exceptions:

- Mail, Scheduling & CalConn MA
- SMTP/MTA MA

These databases were not installed on *YourPrimarySvr* because the associated server tasks do not run on that system.

You now want to install Pinnacle on ***YourSpokeSvr1/YourCompany***.

Since both CalConn and SMTP are to run on *YourSpokeSvr1*, install the above-named databases on that system. Monitor can then efficiently detect and correct issues involving those tasks.

D.1.2 Here's the rub...

When presented with the option of creating IntelliWatch databases on *YourSpokeSvr1* from the media or via replication from another server, you would normally choose the latter option, and name *YourPrimarySvr/YourCompany* as the replica server.

The two MAs that weren't installed on *YourPrimarySvr* cannot be replicated to *YourSpokeSvr1*, however. Therefore, when copies of these databases are not found on *YourPrimarySvr*, the Setup copies them directly from the media.

So far, there is still no issue. *However*, if you have additional servers running those tasks, an issue *will* arise.

Let's say you want *YourSpokeSvr2* also to run CalConn and SMTP (and hence to have copies of the IntelliWatch MAs customized for these tasks). If you replicate IntelliWatch databases from *YourPrimarySvr* during the Setup, these two databases on *YourSpokeSvr2* will have different Replica IDs than the corresponding copies on *YourSpokeSvr1*.

Remember ...

They can't be replicated from YourPrimarySvr because they don't exist there. The Setup, when unable to replicate them, defaults to the copies on the media.

Theoretically, these two databases could be created on *YourSpokeSvr2* (post-setup) by replicating them from *YourSpokeSvr1*. To maintain them, however, would involve using a second replica server for these two MAs, unnecessarily complicating (and perhaps compromising) efforts to maintain database consistency across your environment.

D.1.3 A simple solution

Three simple measures can be taken that eliminate issues like those in the above example:

- install all Pinnacle components on the Primary Server
  - This applies even if certain of the associated add-in tasks will not be running there.

    This guarantees that you have a Master Set of *all* IntelliWatch databases that can be replicated to other systems on subsequent installs.
- replicate IntelliWatch databases from this *one* server only
  - This ensures that all IntelliWatch databases that *do* replicate will have the same Replica ID.
- do all editing of databases on the Master Set
  - This replication can be set up as PUSH-only, preventing changes made in databases on Spoke Servers from contaminating the Master Set.

The first of these recommendations is not self-evident, since installing software components on a system where they will not be used should usually be avoided. Nevertheless, the advantages gained as regards the relative simplicity (and consistency) of your replication regimen for IntelliWatch databases makes it well worth implementing.

Appendix

NT Setup Dialogs

E

This discussion of Setup dialogs is supplemental to the *"Installation Guide" on page 31*.

Full installation:

Please note that the following dialogs are in the order in which they are displayed during a full installation. Depending on your installation choices, however, the Setup may skip certain dialogs.

Upgrade:

Dialogs that are displayed only during an upgrade are discussed after the main series. For all dialogs in the upgrade path, start with *E-1* through *E-6*, then follow the ***Upgrade Path*** links toward the bottom of the respective pages. (See also *"Upgrade Applications" on page 439*.)

Cancel and Back buttons:

The Back button behaves as expected, and is not discussed below.

The Cancel button has two behaviors: On a main dialog, it allows you to exit (or resume) the installation. On a sub-dialog (launched by the Browse button, for example) it returns you to the main series of Setup dialogs.

FIGURE E-1: Location to Save Files

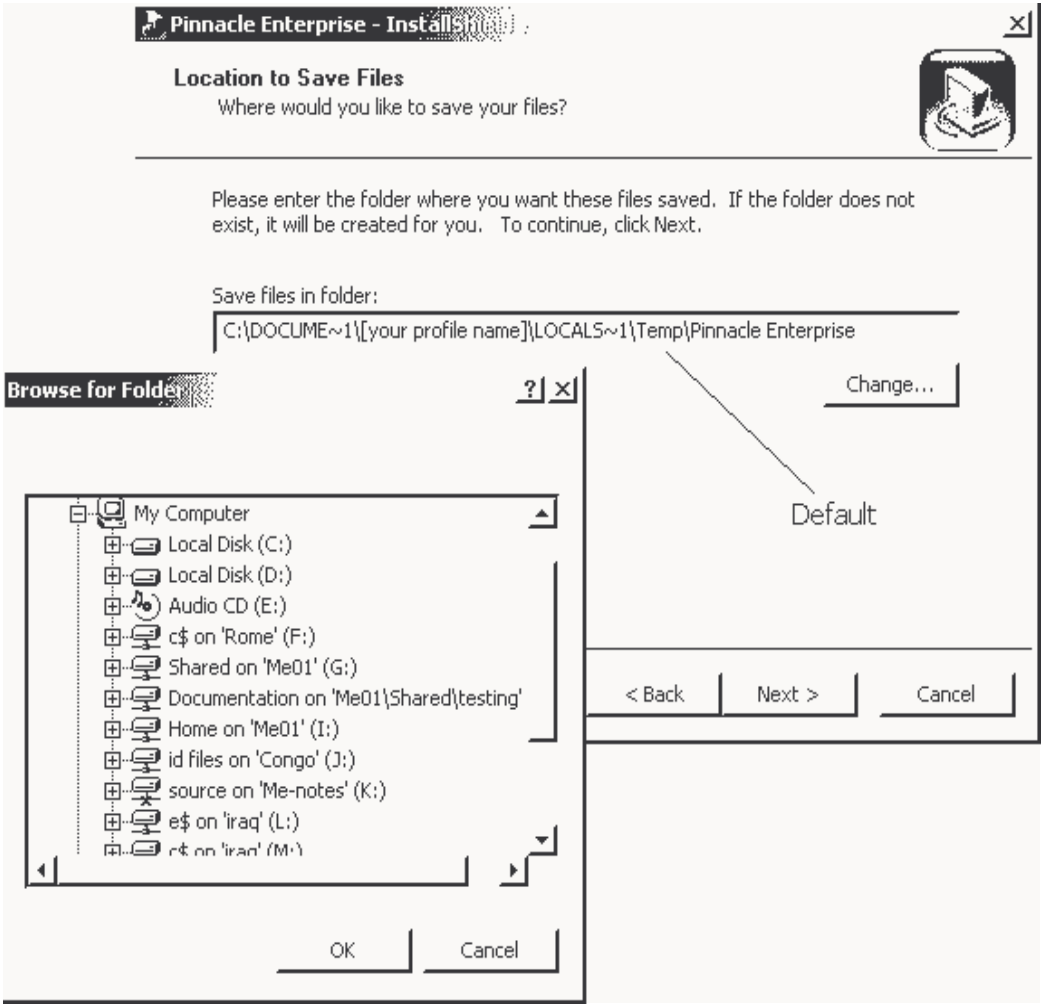

Purpose of dialog:

Allows you to accept the default file-extraction location, or to customize it.

Click Next to accept the default.

Click Change to modify the location. This brings up the dialog in the foreground. Use this dialog to navigate to the desired location, then click OK.

The new location appears in the text box. Click Next to proceed.

FIGURE E-2: Extracting Files

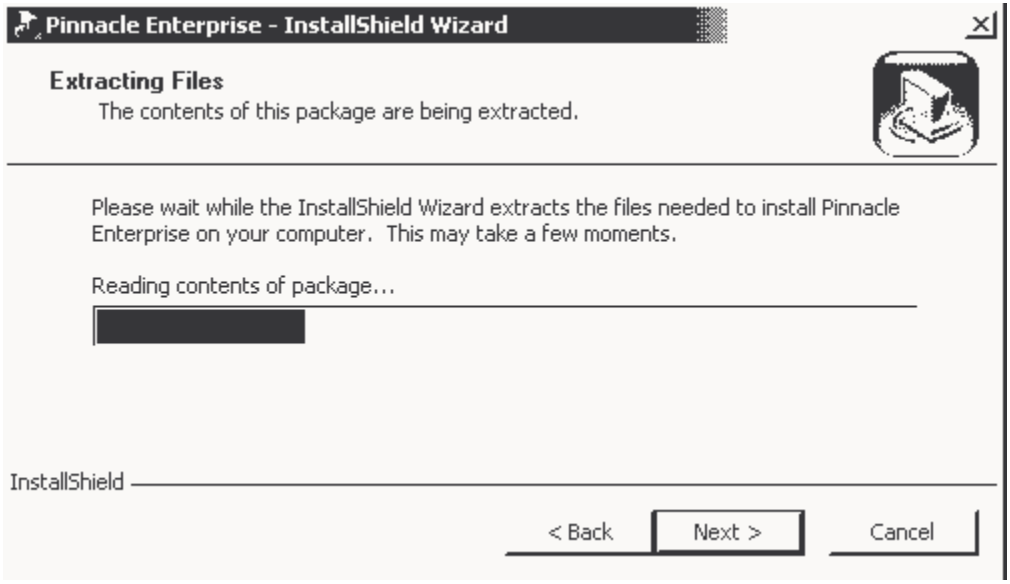

Purpose of dialog:

Allow you to follow the progress of file extraction.

When file extraction has completed, the dialog on the following page is displayed.

Cancel is the only button enabled while files are extracting.

FIGURE E-3: Overwrite Protection

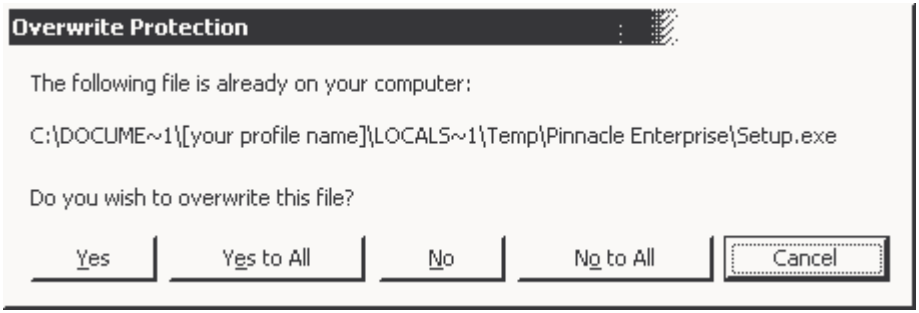

Purpose of dialog:

The Overwrite Protection dialog appears only if a version of the IntelliWatch Setup is detected on your system.

The dialog buttons function as follows:

- Yes
  - Click this button to overwrite--one-by-one--Pinnacle Enterprise Setup files already on the system.
- Yes to All
  - Click this button to overwrite--all at once--Pinnacle Enterprise Setup files already on the system.
- No
  - Click this button to keep--one-by-one--Pinnacle Enterprise Setup files already on the system.
- No to All
  - Click this button to keep all Pinnacle Enterprise Setup files already on the system.

If you selected any of the above four buttons, the Setup will proceed with the dialog at *Figure E-4* once all files have been processed.

FIGURE E-4: Choose Setup Language

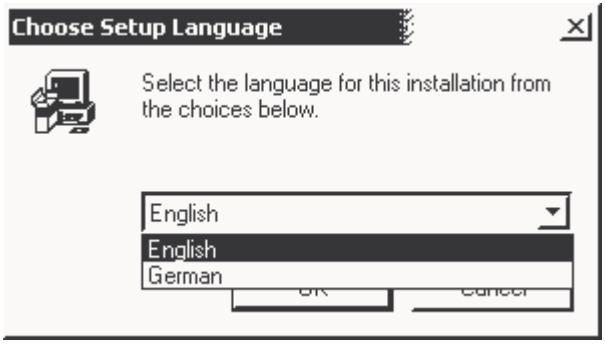

Purpose of dialog:

Currently, there are two options: English (the default) and German:

- English
  - The standard, all-English-language version of the product is installed.
- German
  - The standard, all-English-language version of the product is installed, with two exceptions:
    - most text on the Setup dialogs is displayed in German
    - search strings of Database Scan Triggers are in German

All other features of the installed are in English, including IntelliWatch databases, on-screen help texts, menu items, and so on.

Once you have made your selection, click OK to proceed.

FIGURE E-5: IntelliWatch Setup

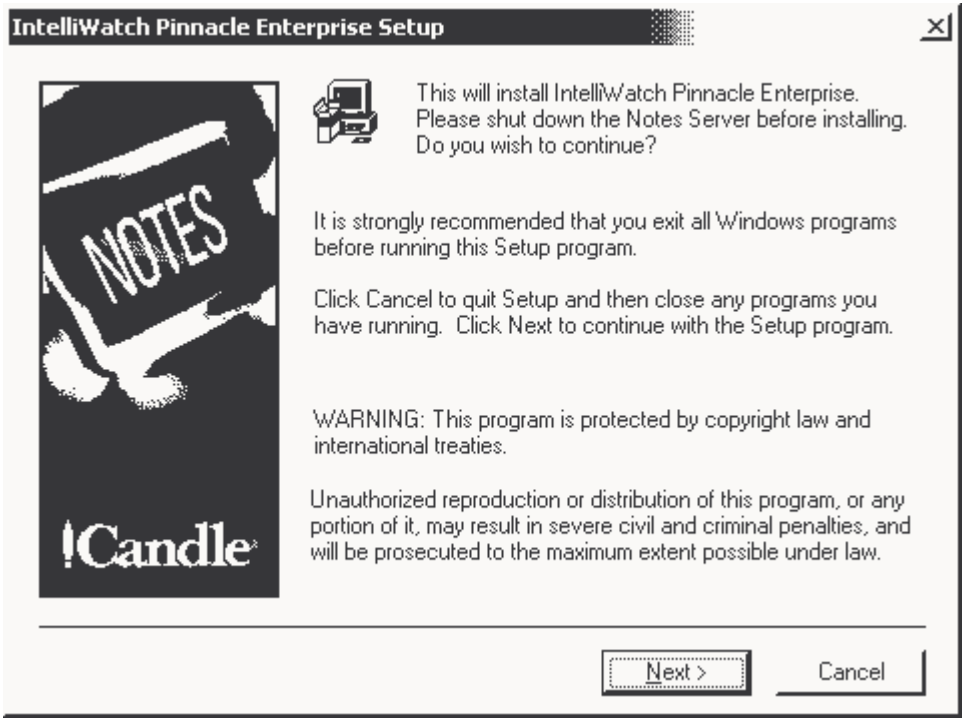

Purpose of dialog:

To advise you to do the following before proceeding with the Setup:

- shut down the Notes server
- close all Windows programs

To remind you to respect all applicable copyright laws when installing or distributing this program.

Click Next to proceed.

FIGURE E-6: IntelliWatch Pinnacle License Agreement

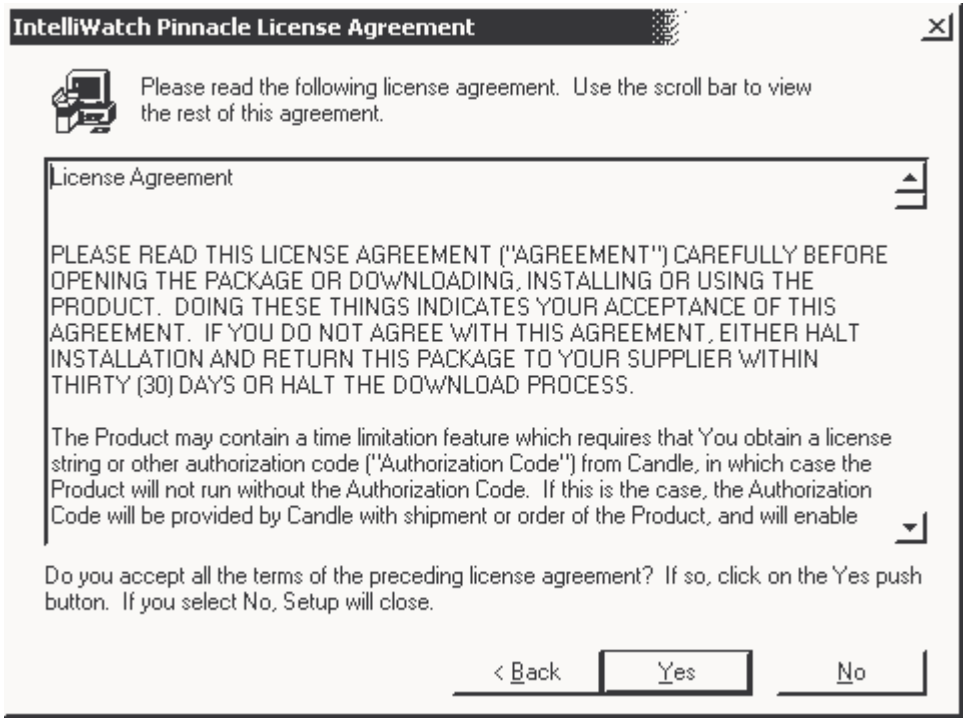

Purpose of dialog:

To display the license agreement.

Click Yes to accept all terms of the displayed license agreement and proceed to the next dialog.

Click No if you do not accept these terms. The Setup will close.

Upgrade path:

Assuming all Pinnacle 99 components were installed on the system in question, go to *E-10*.

FIGURE E-7: IntelliWatch Setup Options

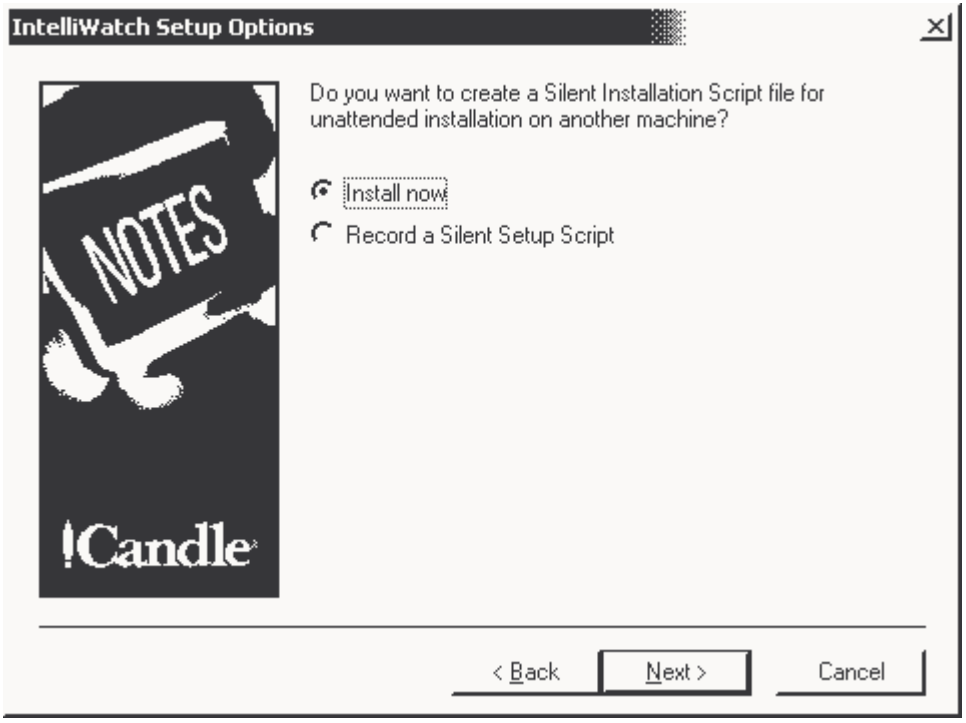

Purpose of dialog:

Gives you the option of installing the product immediately, or recording a Silent Setup Script to be used for future installations in your environment.

If you elect to record a Silent Setup, your choices are recorded in a file, **iwsetup.ini**. For more information on this latter option, see *"Silent Setup" on page 38*.

Make your selection and click Next to proceed.

FIGURE E-8: IntelliWatch Setup.INI File

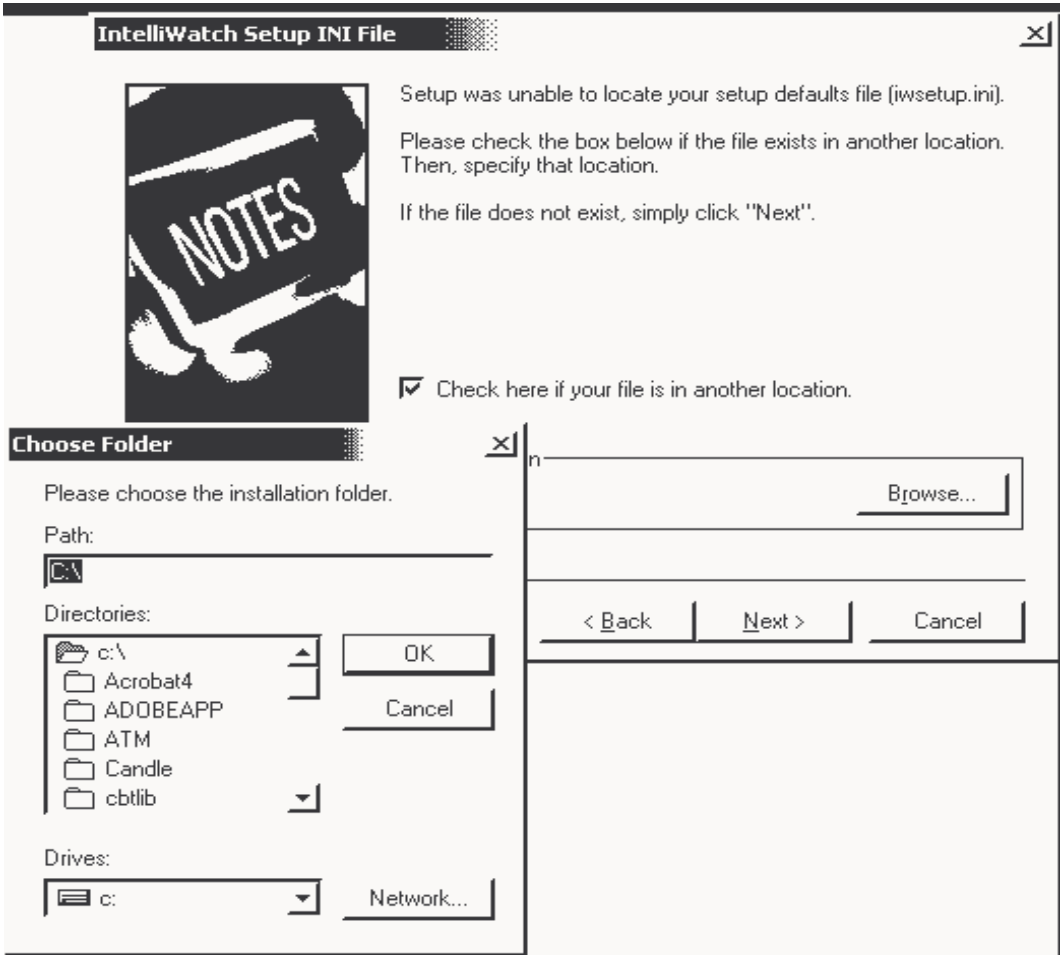

Purpose of dialog:

To allow you to select a Silent Setup Script (**iwsetup.ini**) as the source of IntelliWatch configuration settings.

The Setup looks for this file in the directory from which Setup.exe was launched. To access an **iwsetup.ini** file in a different location, select the checkbox in the middle of the dialog, then click Browse to bring up the foreground dialog. Navigate to the desired location, then click OK.

Finish this step of the Setup by clicking Next.

FIGURE E-9: IntelliWatch Setup Types

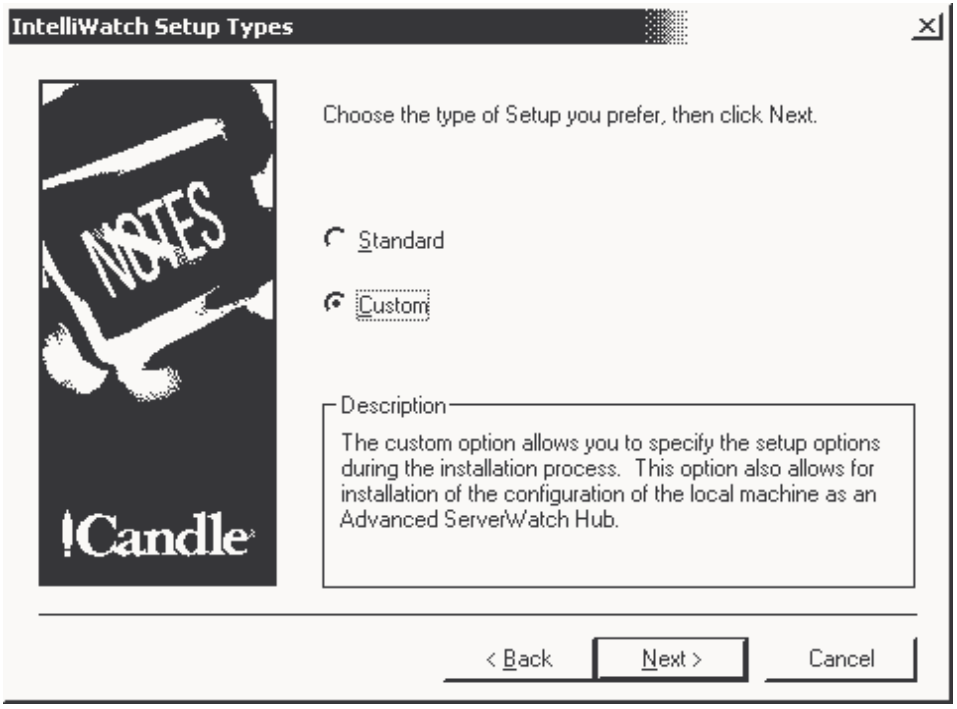

Selecting Standard causes the following Pinnacle server components to be installed:

- Monitor
- Performance Manager
- Tracer

Among the options disabled by default are the following:

- Monitor Availability Stats
- Tivoli TEC Event (Command option)
- NT Event Log (Command option)
- sending Paging (from the UI)
- sending SNMP traps (from the UI)

Custom must be selected to install the Advanced ServerWatch Hub task.

A ***Standard*** installation continues with dialogs: F-13-16; F-22-23; F-33-37.

FIGURE E-10: Active NT Services

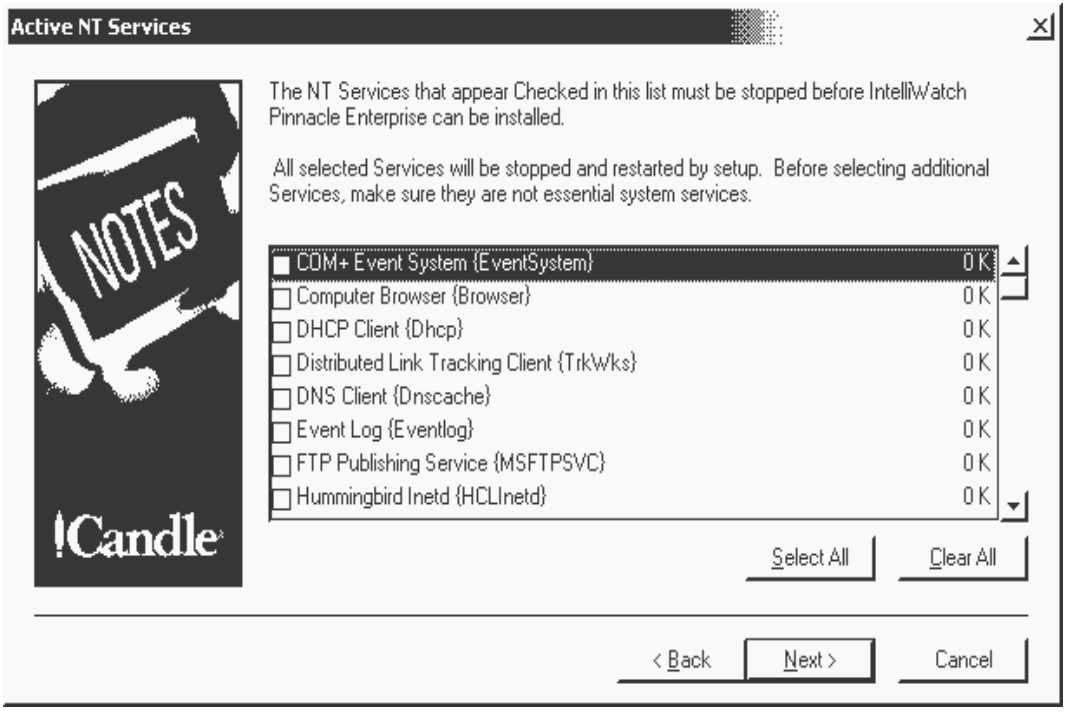

Purpose of dialog:

The Setup detects NT Services installed on the system, and stops and restarts them as needed.

Should you have an NT Service installed on your system that is not detected by the Setup, stop it manually, before continuing.

When you have confirmed that the appropriate Services have been selected, click Next to proceed.

Upgrade path:

Assuming all Pinnacle 99 components were installed on the system in question, go to *E-40*.

FIGURE E-11: Select IntelliWatch Components

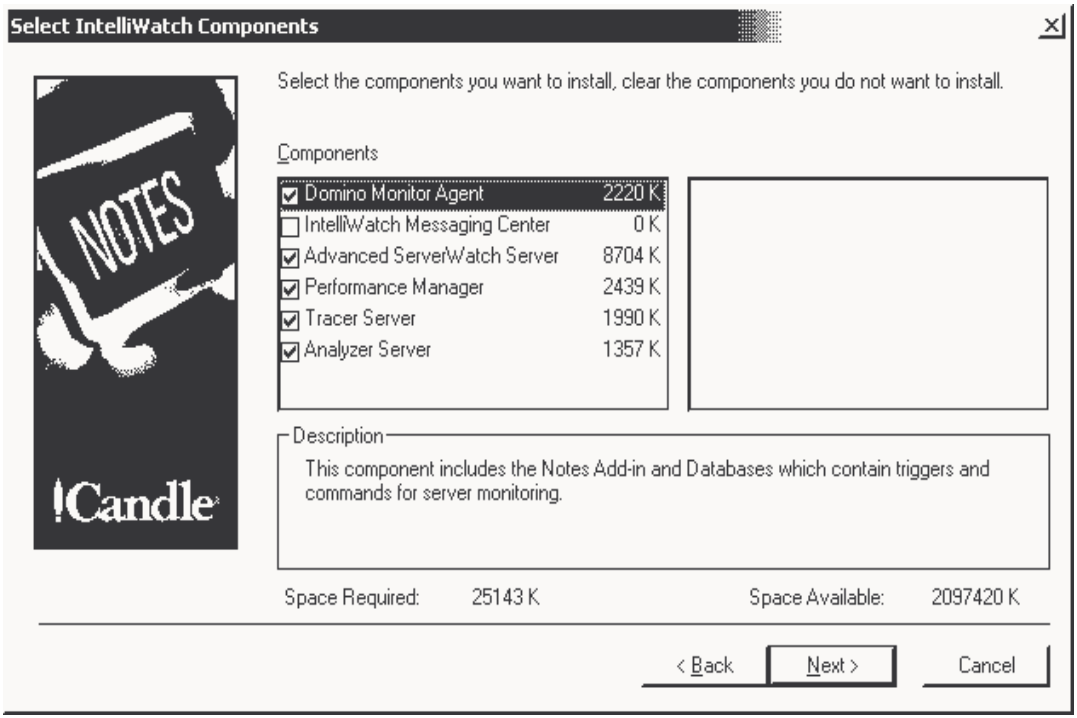

Purpose of dialog:

To allow the selection of individual IntelliWatch components.

If you selected Standard as the Setup Type (at *Figure E-9*), this dialog is not displayed.

A medium-sized Notes environment (30 servers) usually runs only one instance of Analyzer Server, two instances of Advanced ServerWatch and one IntelliWatch Messaging Center.

For details, see "Sample architecture" on page 32.

Once you have made your Pinnacle component selections, click Next to proceed.

Upgrade path:

Assuming all Pinnacle 99 components were installed on the system in question, go to *E-41*.

FIGURE E-12: Installation Server Type

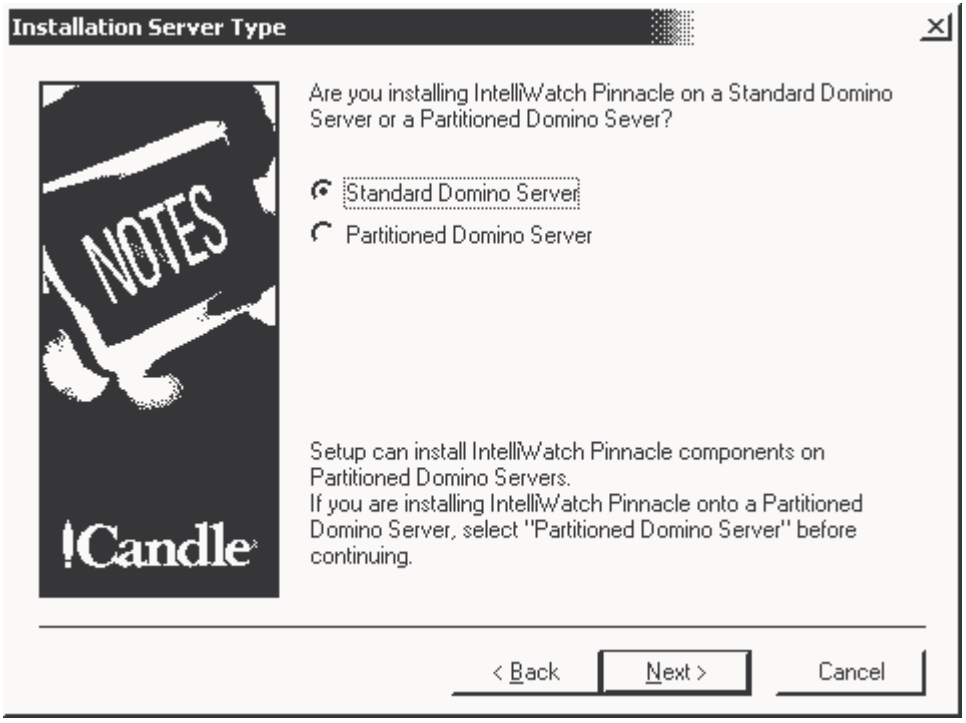

Purpose of dialog:

To allow you to select between server types. On Domino 5.0, the Setup checks for Notes partitioning in the system registry (under Lotus); on R4, the Setup looks in the Notes data directory for a batch file containing the string ***Set NOTESPARTITION=[number]***.

If this server is partitioned, yet this is not detected by the Pinnacle Setup--or vice versa--we suggest you verify your system configuration before proceeding.

*Occasionally, the **nserve.bat** file remains on a system that formerly was partitioned at the Notes level, but now is running a standard Domino server. (The contents of this file set the Notes partition number.) If you encounter such a file on a standard Domino server, verify your Notes configuration before proceeding.*

Once the server type is confirmed, click Next to proceed.

FIGURE E-13: Choose Location of Install Directories

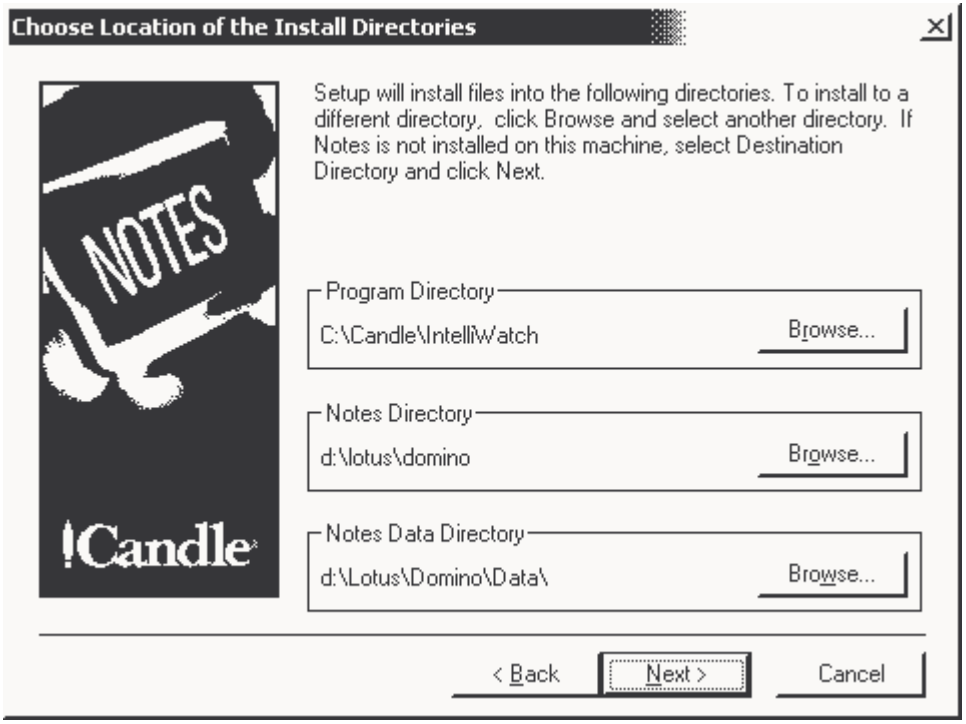

Purpose of dialog:

The Candle directory displayed above is the default. Do nothing to keep the default. Click Browse if you want to customize this directory.

The Notes directories displayed above were detected by the Setup. If these locations are not correct, change the location(s) by means of the dialogs brought up by clicking on the Browse buttons.

Once the desired install directories are displayed, click Next to proceed.

The executables that launch the IntelliWatch server tasks are located in the Notes directory.

FIGURE E-14: Choose Destination Location

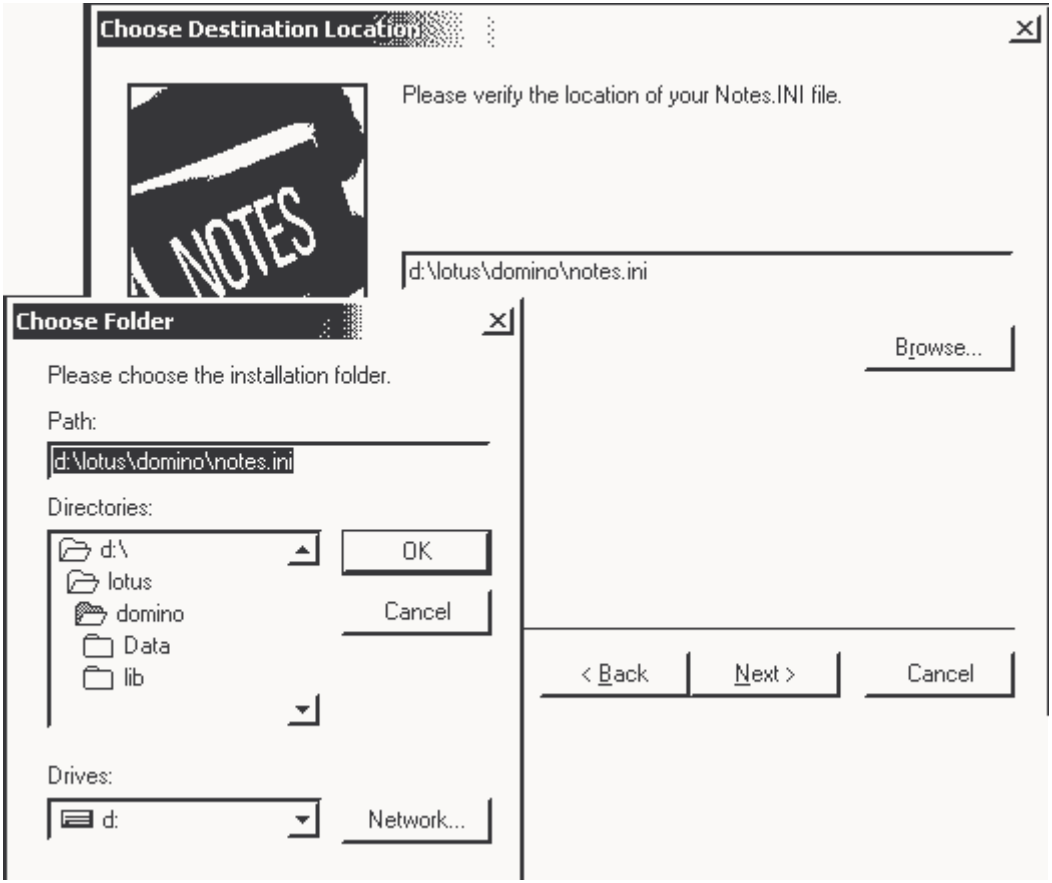

Purpose of dialog:

To allow you to verify the location of the **notes.ini** detected by the Setup.

To select a **notes.ini** other than the one detected by the Setup, click the Browse button, which brings up the foreground dialog. Navigate to the file, then click OK to return to the main dialog.

When the correct location is displayed in the text box of the original dialog.

Once the location of the **notes.ini** has been verified, click Next to proceed.

Upgrade path:

Assuming all Pinnacle 99 components were installed on the system in question, go to *E-17*.

FIGURE E-15: IntelliWatch Pinnacle Registration

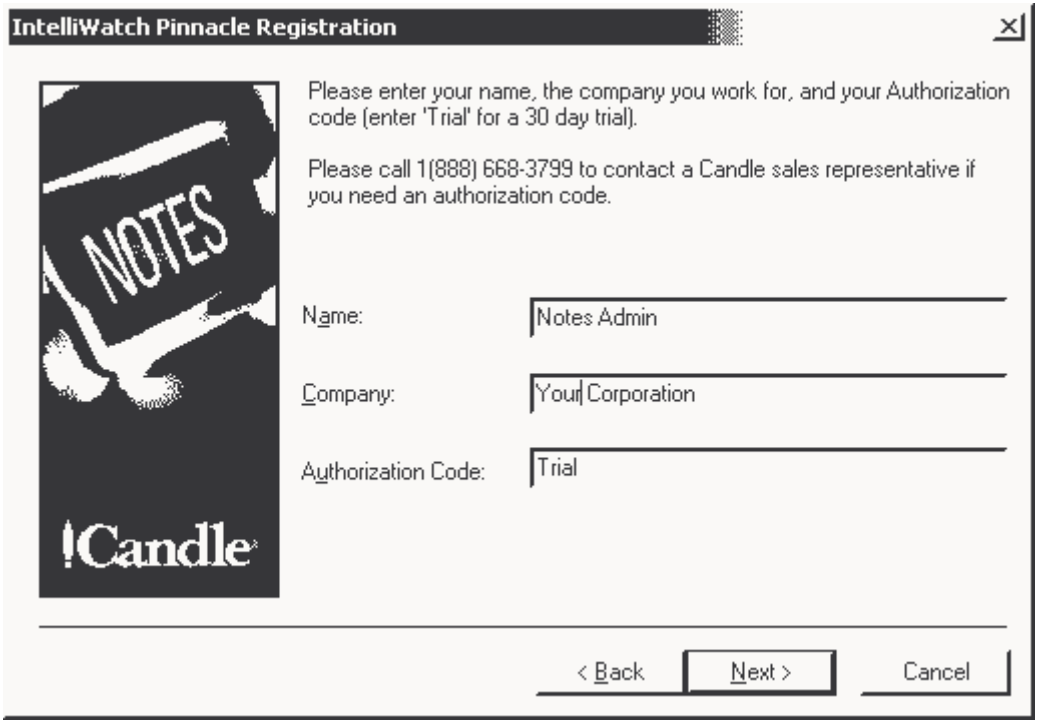

Purpose of dialog:

To allow you to enter a the name of the relevant Admin, as well as the name of the company.

If you have purchased the product, enter the authorization code you received in the bottom text box. If this is a product trial, enter "Trial" (without quotes) in the textbox.

Once you have entered all registration information, click Next to proceed.

The Next button remains disabled until all three fields contain a value.

This and the following dialogs will not be displayed with maintenance release 27.37.

FIGURE E-16: IntelliWatch Pinnacle Authorization Code

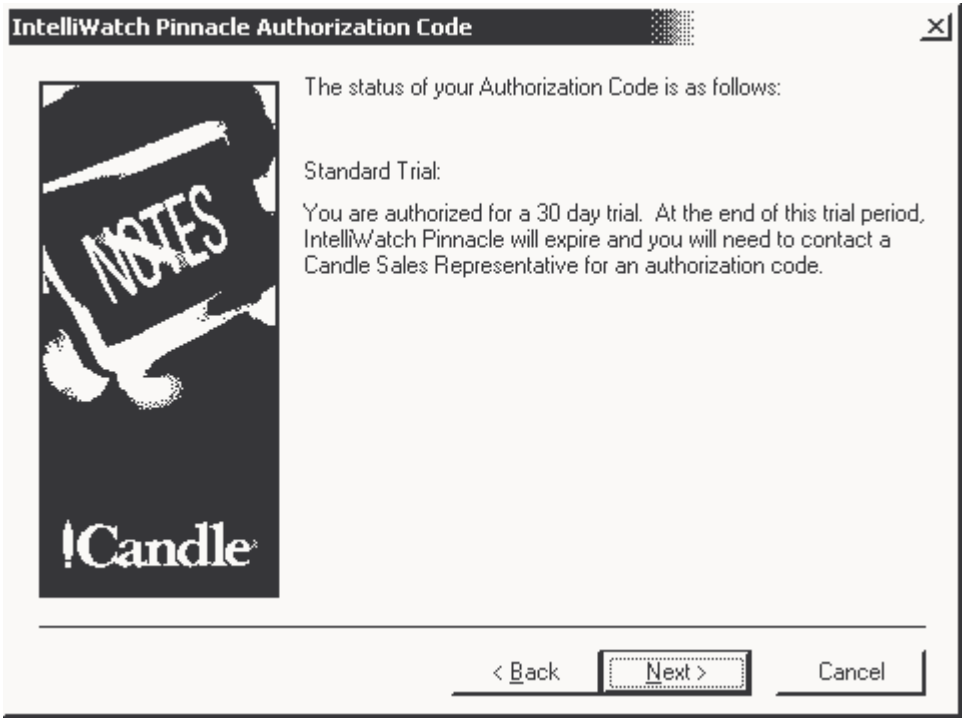

Purpose of dialog:

To allow you to verify your authorization code.

If you made an error on the previous dialog, click Back to correct it. Otherwise, click Next to proceed.

This and the previous dialogs will not be displayed with maintenance release 27.37.

FIGURE E-17: Database Target Subdirectory

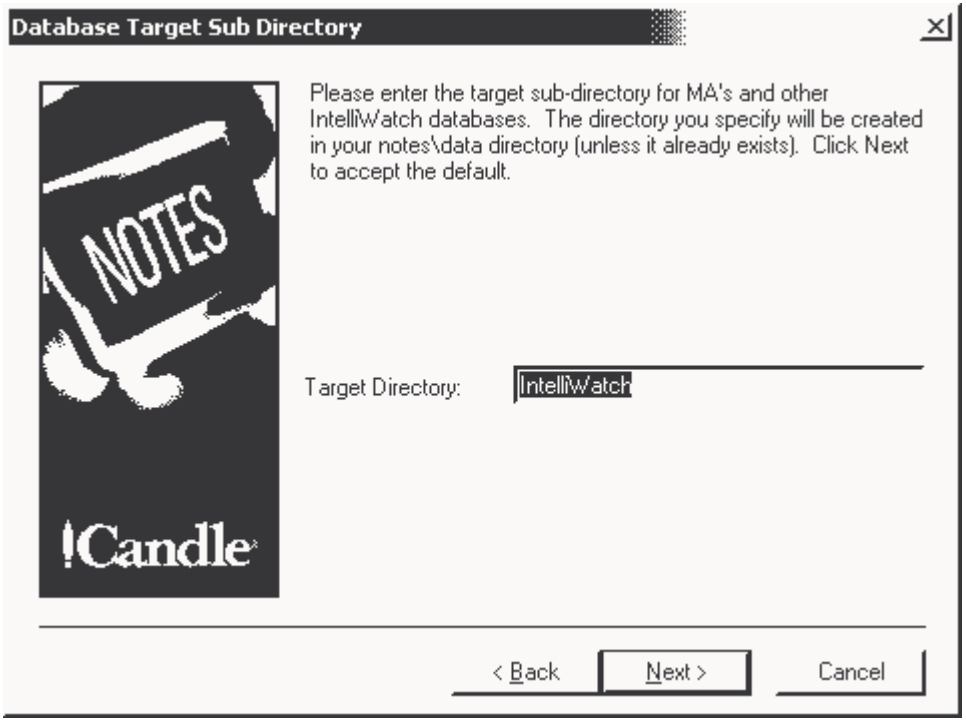

Purpose of dialog:

To allow you to customize the name of the subdirectory that contains IntelliWatch databases.

The directory will be created beneath your Notes data directory. If that latter directory were called **D:\NotesData**, for example, and you entered **OurIWDbs** as a target directory, IntelliWatch databases would be located in **D:\NotesData\OurIWDbs**.

Once you have made any desired changes in the name of the target directory, or simply to accept the default, click Next.

Upgrade path:

Assuming all Pinnacle 99 components were installed on the system in question, go to *E-18*.

FIGURE E-18: Database Source Selection

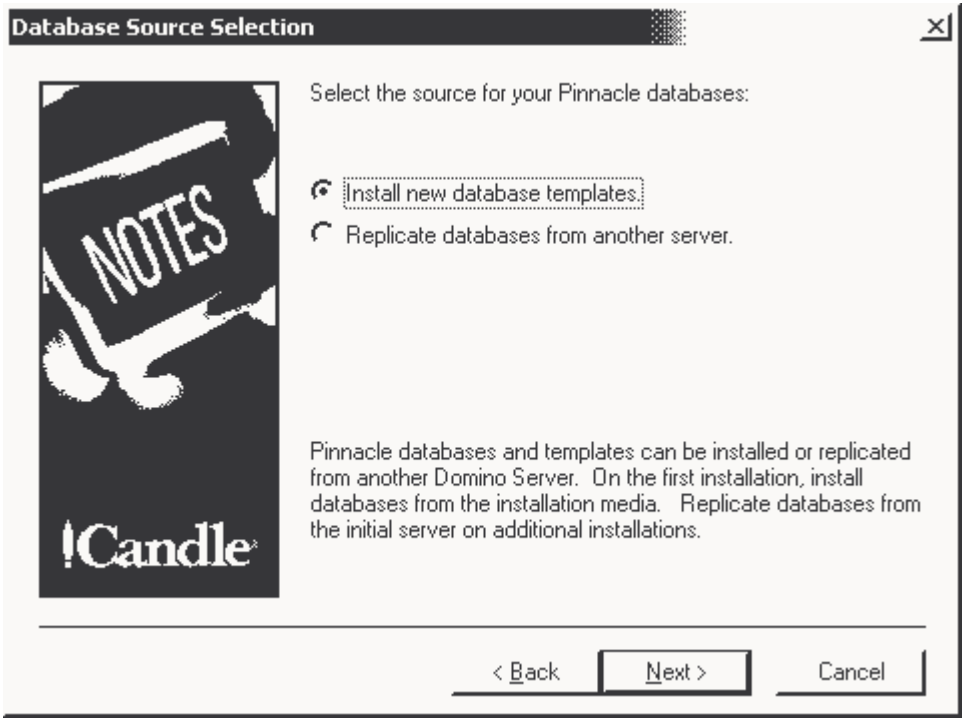

Purpose of dialog:

Use to select the source of IntelliWatch Pinnacle databases for installation on this system.

Install IntelliWatch databases the media on the first server only, then replicate them from that server to the system on which IntelliWatch is currently being installed. (See also "Installation guidelines" on page 32, and "MA Deployment" on page 45.)

For specifics regarding IntelliWatch databases and replication, see "IntelliWatch Databases: Template Usage and Replication" on page 395.

Once you have selected the source of IntelliWatch databases, click Next to proceed.

Upgrade path:

Assuming all Pinnacle 99 components were installed on the system in question, go to *E-24*.

FIGURE E-19: IntelliWatch Trigger Selection

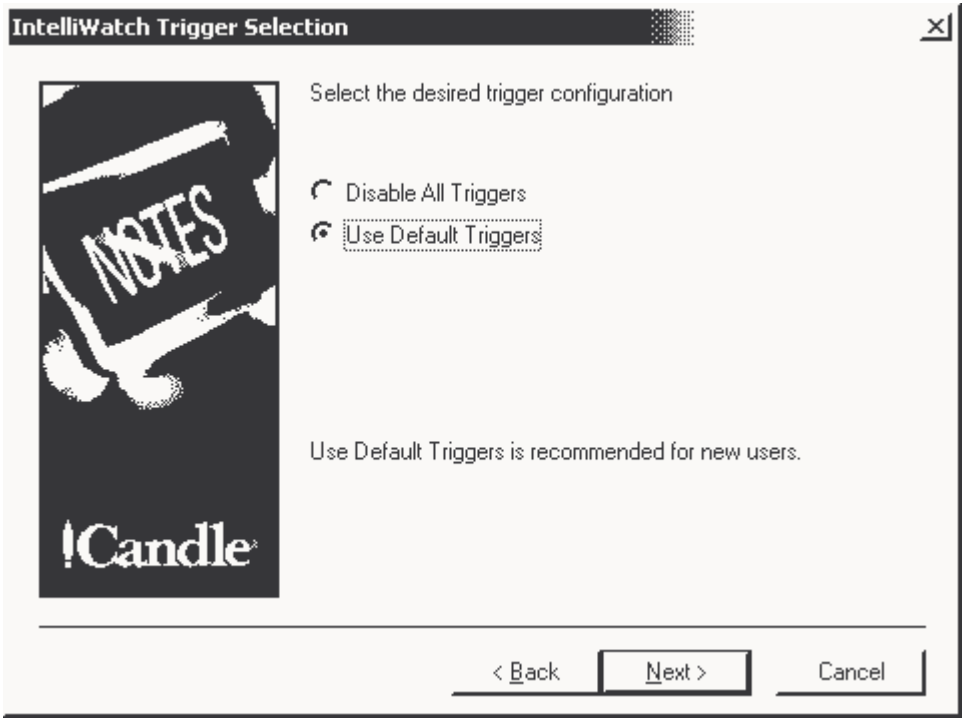

Purpose of dialog:

To allow you to disable all Triggers, or enable selected Triggers during the installation process.

Select "Disable All Triggers" when installing on production systems. Crash Detection will be active (see page 279), but no Triggers will be monitoring server activity. Activate Triggers on your Primary Server, then replicate MAs as per your system architecture.

Select "Use Default Triggers" when doing a Trial installation on test systems. A number of Triggers that have proven useful in many environments will be activated, and you will be able to see firsthand how IntelliWatch detects and corrects server issues.

Once you have selected the desired Trigger configuration, click Next to proceed.

FIGURE E-20: Select Replication Server

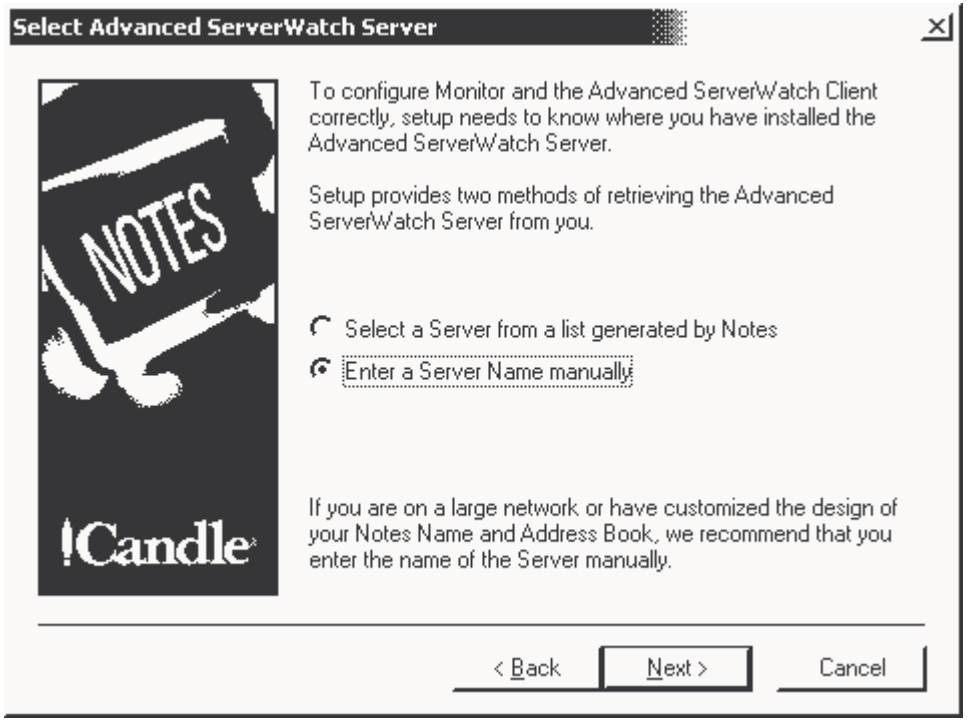

Purpose of dialog:

To allow you flexibility in selecting the Replication Server.

As noted on the dialog, above, if the design of your Domino Directory has been customized, or if your network is particularly large, we suggest you enter the server name manually.

Once you have made your selection, click Next to proceed.

FIGURE E-21: Select Replication Server

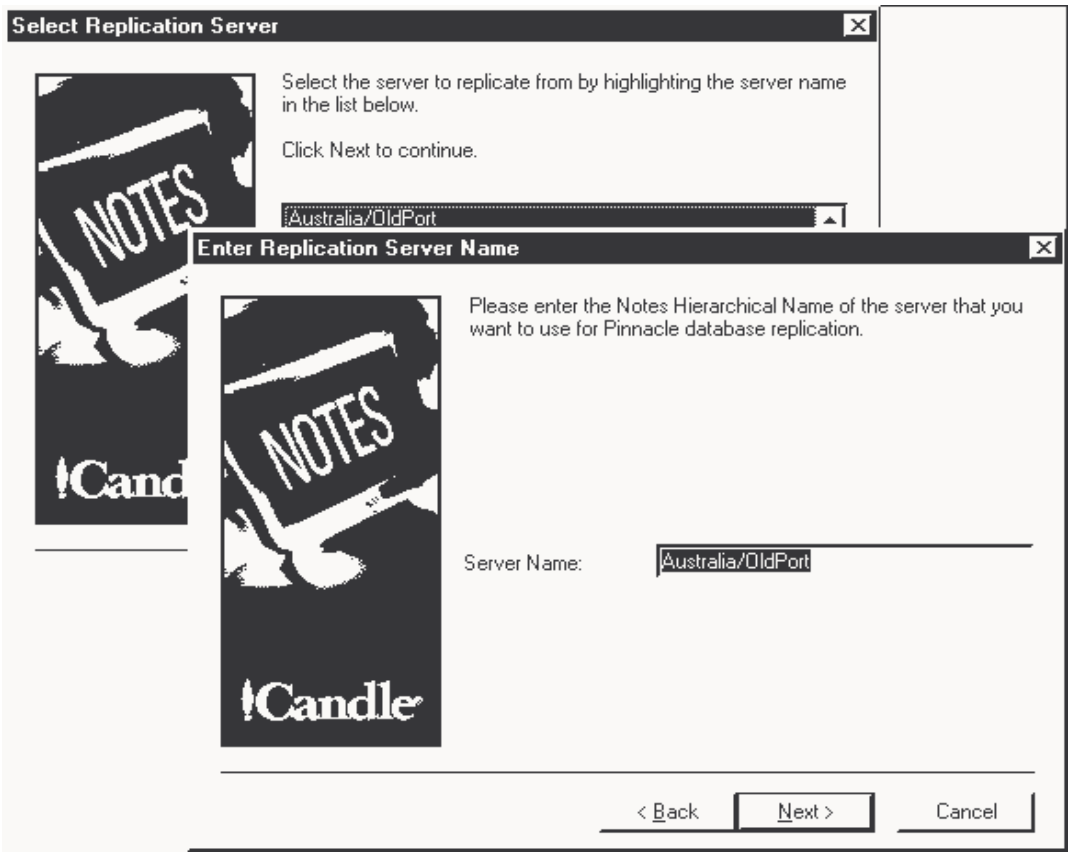

Purpose of dialog:

To allow you either to select the Replication Server from a list generated by Notes, or to enter the server name manually.

If you selected the first option, the dialog in the background, above, is displayed. Click on the server to select it.

If you opted to enter the server name manually, the dialog in the foreground is displayed. Enter the server name in the text box.

In either case, click Next to proceed.

FIGURE E-22: Replication Source Subdirectory

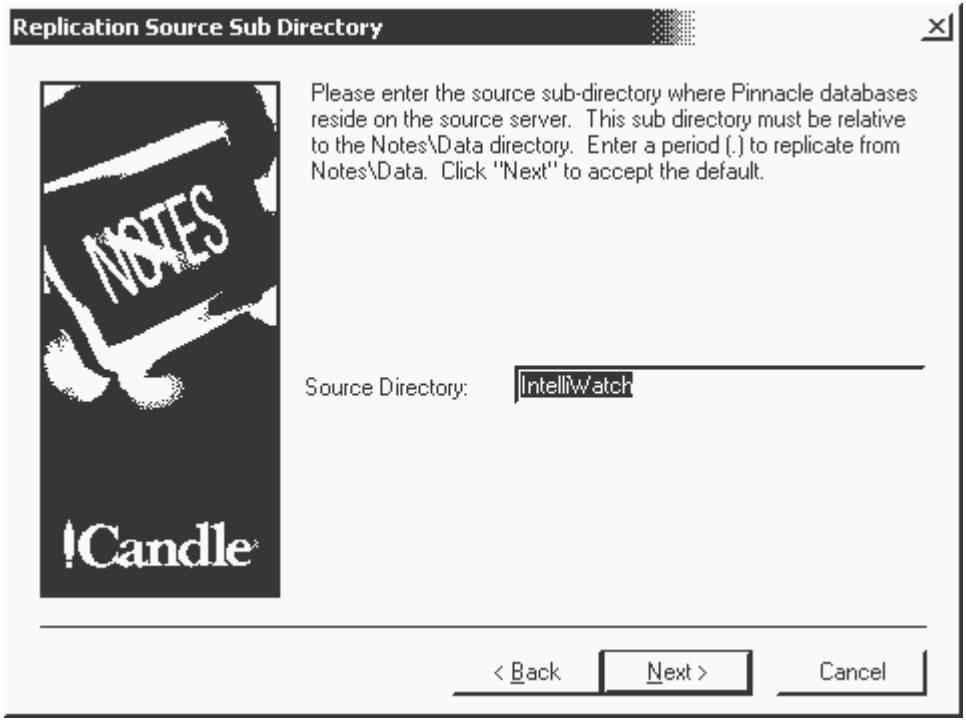

Purpose of dialog:

To allow you to customize the subdirectory where the Setup should look for IntelliWatch databases that are to be replicated.

Normally, this will be the same as the location specified at *Figure E-17*, and no editing should be needed.

Once the text box contains the correct source directory, click Next to proceed.

FIGURE E-23: IntelliWatch Pinnacle Network Parameters (Server)

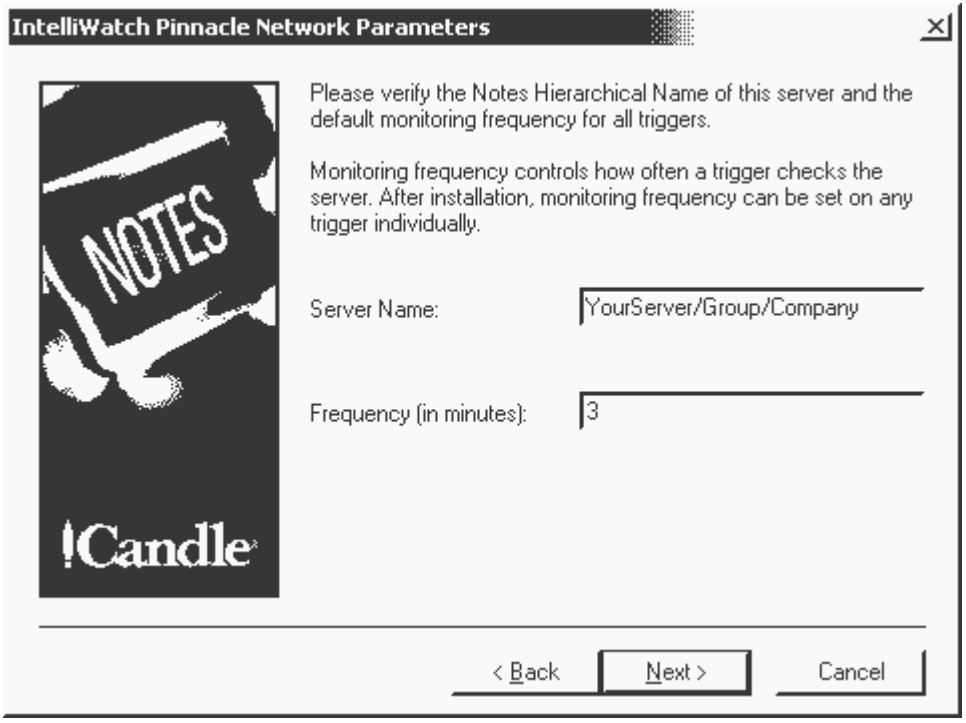

Purpose of dialog:

To allow you to verify the name of the Domino server on this system, and to customize the monitoring frequency for IntelliWatch Triggers (the default is 15 minutes). (Under normal circumstances, the Server Name displayed should be correct. If it isn't, please verify your Notes configuration.)

Consider carefully the ramifications of the Frequency parameter. Every active Trigger uses this default frequency, unless a different interval is specified for that Trigger.

Also, bear in mind that certain Trigger types (such as Replication Integrity, for instance) use a considerable amount of system resources when they evaluate. You may wish to adjust either the frequency or the time of day at which they evaluate--or both.

Once the text boxes contain the correct information, click Next to proceed.

FIGURE E-24: IntelliWatch Pinnacle Network Parameters (Admin)

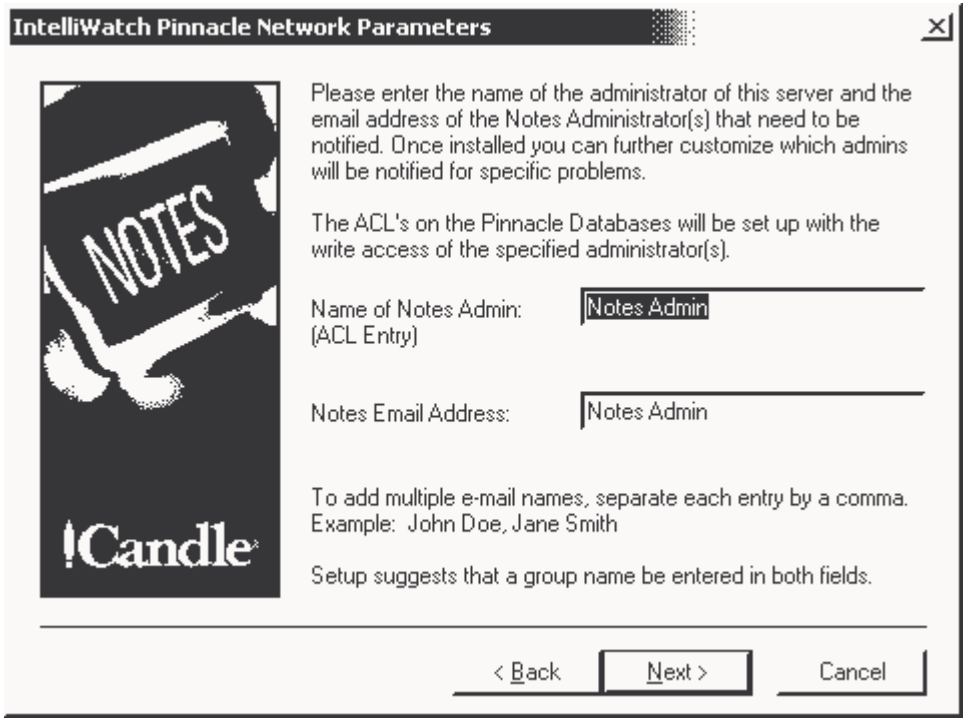

Purpose of dialog:

To allow you to specify the name and e-mail address of the Notes Admin(s) who will have administrative authority over this server.

Note that the top text box should contain the name of the Admin(s) as you want them to be created in the ACL of IntelliWatch databases.

Separate multiple entries by commas. Use of groups rather than individual names is recommended.

Once the text boxes contain the correct information, click Next to proceed.

Upgrade path:

Assuming all Pinnacle 99 components were installed on the system in question, go to *E-25*.

FIGURE E-25: Notification Mechanisms

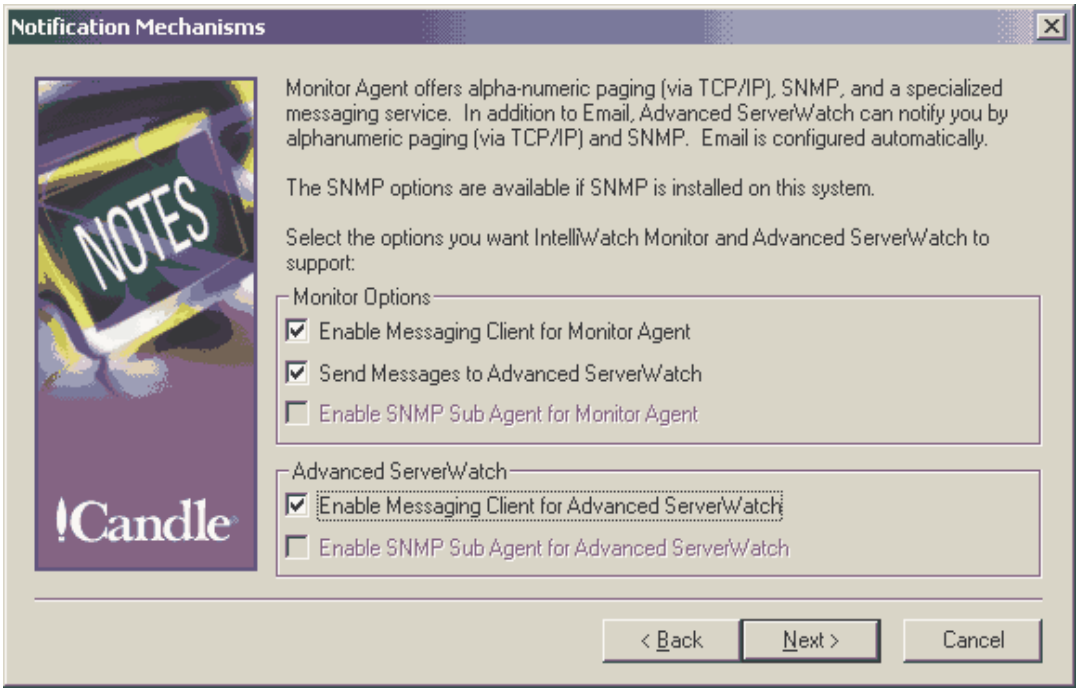

Purpose of dialog:

To allow you to customize the notification mechanisms used by Monitor and Advanced ServerWatch.

While e-mail notification is configured automatically, alphanumeric paging and SNMP traps from Monitor and Advanced ServerWatch must be selected here, or they will not be available.

If the appropriate protocols/resources are not available on (or from) a system, certain options will be unavailable. SNMP was not installed on the machine running the Setup for this dialog, for example, and SNMP options are therefore grayed out.

Once you have made your selections, click Next to proceed.

Upgrade path:

Assuming all Pinnacle 99 components were installed on the system in question, go to *E-28*.

FIGURE E-26: IntelliWatch Analyzer Options

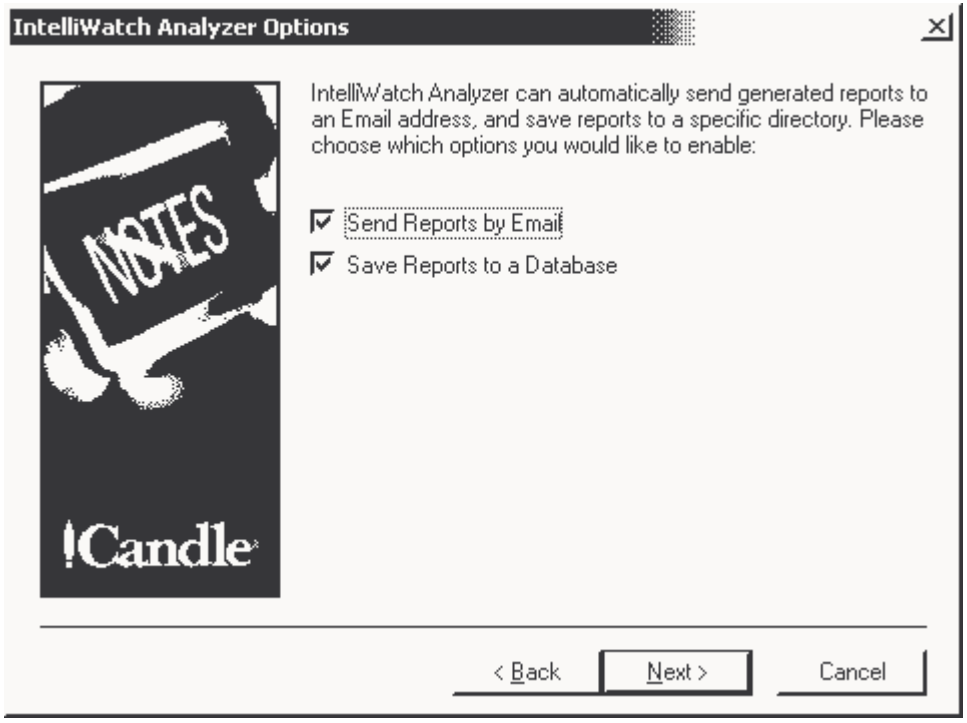

Purpose of dialog:

To allow you to send Analyzer reports via e-mail, as well as save them to a database.

E-mail will be sent to the address(es) listed in the NT registry under HKEY_LOCAL_MACHINE\SOFTWARE\Candle\IntelliWatch\Analyzer\Preferences\Send To. (The registry key in question is populated by the Notes E-mail Address field on dialog *E-24*.) This parameter may be a comma-delimited list.

Reports will be saved to the database listed in the NT registry under HKEY_LOCAL_MACHINE\SOFTWARE\Candle\IntelliWatch\Analyzer\Preferences\Report Storage Database.

Once you have made your choices, click Next to proceed.

FIGURE E-27: Analyzer Report Options

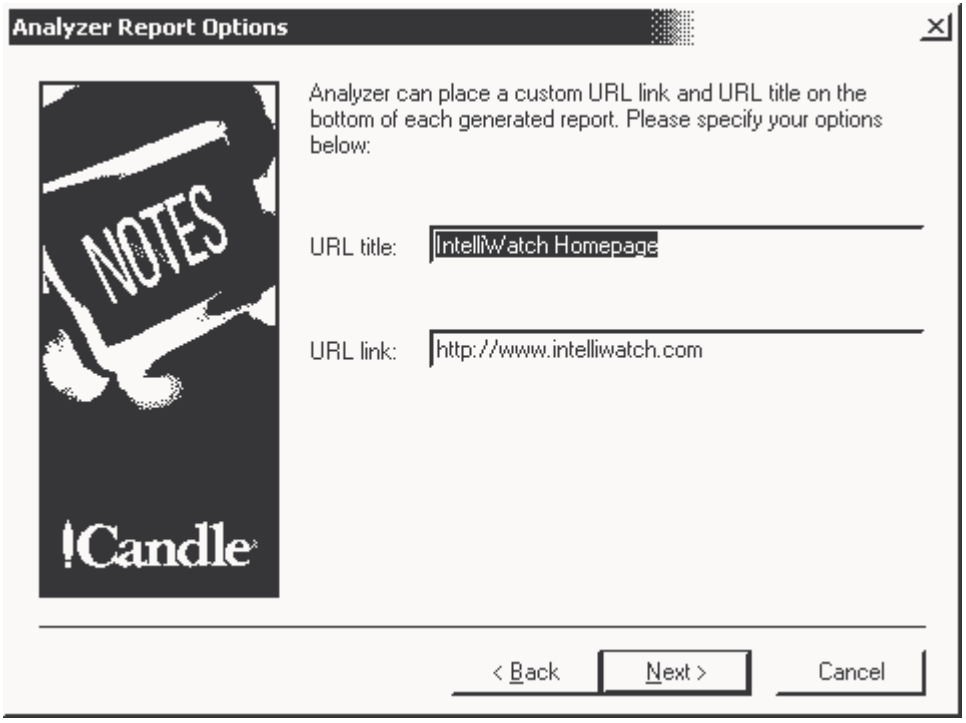

Purpose of dialog:

To allow you to place a custom URL title and/or link at the bottom of each report generated by Analyzer. Both fields are optional.

The URL Title can be whatever you choose.

For the link to work properly, the address must be correct, and the server where the reports are located must be running HTTP.

Once you have made any desired entries, click Next to proceed.

FIGURE E-28: IntelliWatch Messaging Client

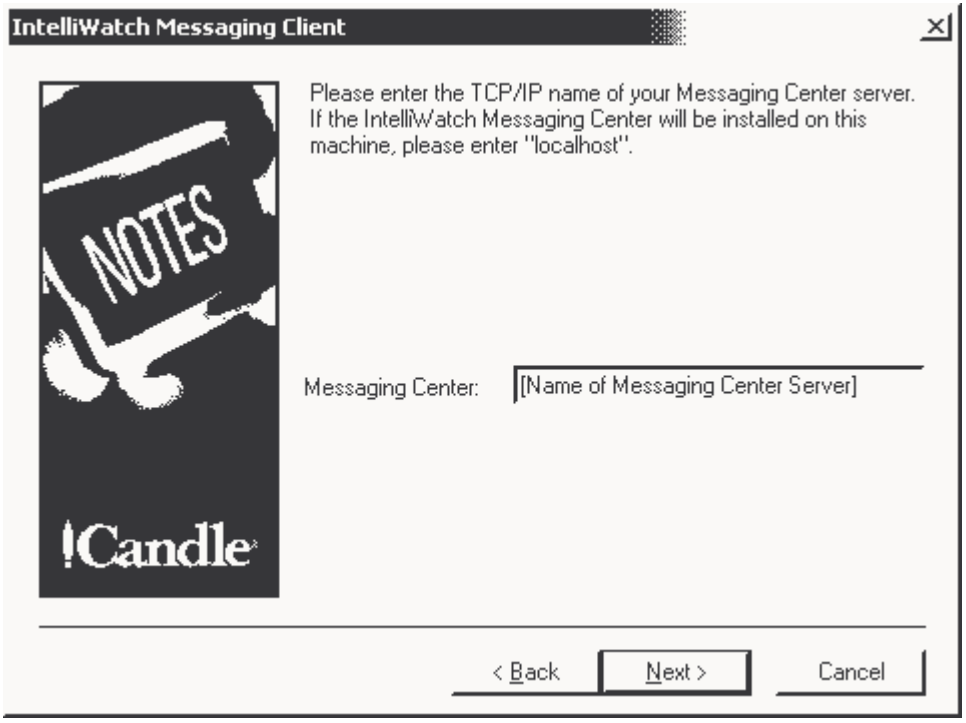

Purpose of dialog:

To allow you to specify the system running the IntelliWatch Messaging Center Gateway. (This service is available on both NT and AIX. For details, see *"IntelliWatch Messaging Center Gateway" on page 273*.)

The IntelliWatch Messaging Client requires that the Messaging Center Gateway be installed on a system accessible (over TCP/IP) from the machine on which the Setup is running.

Once you entered the TCP/IP name of the system running the Messaging Center, click Next to proceed.

Upgrade path:

Assuming all Pinnacle 99 components were installed on the system in question, go to *E-29*.

FIGURE E-29: Messaging Center Options

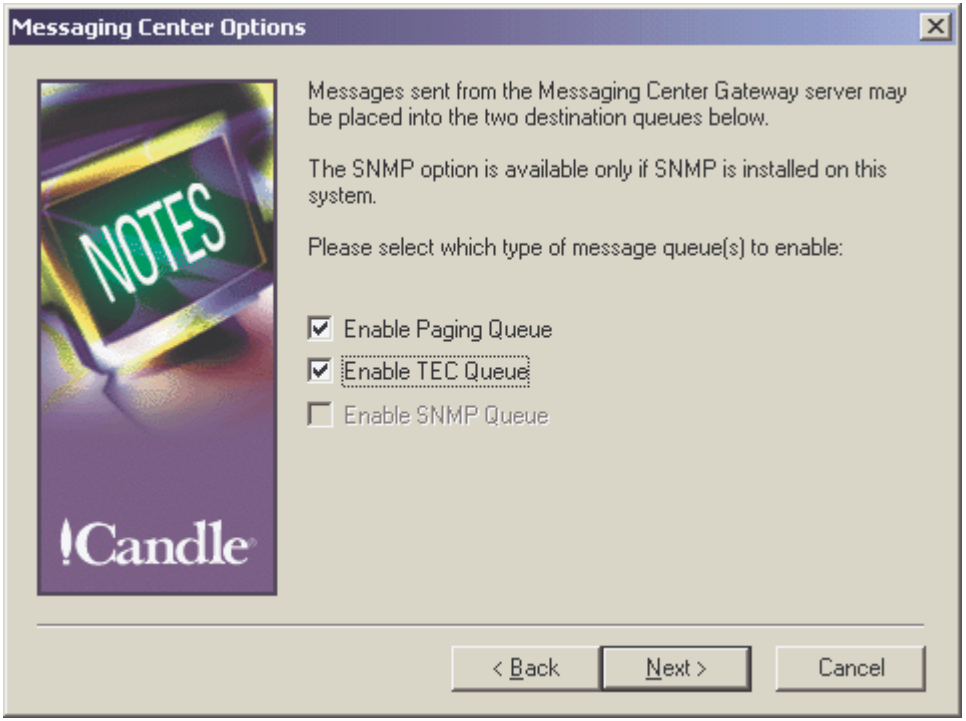

Purpose of dialog:

To allow you to enable queues for SNMP, alphanumeric paging, and Tivoli TEC Events.

If the appropriate protocols/resources are not available on (or from) a system, certain options will be unavailable. SNMP was not installed on the machine running the Setup for this dialog, for example, and the SNMP option is therefore grayed out.

Once you have made any desired entries, click Next to proceed.

Upgrade path:

Assuming all Pinnacle 99 components were installed on the system in question, go to *E-30*.

FIGURE E-30: IntelliWatch Messaging Center Information

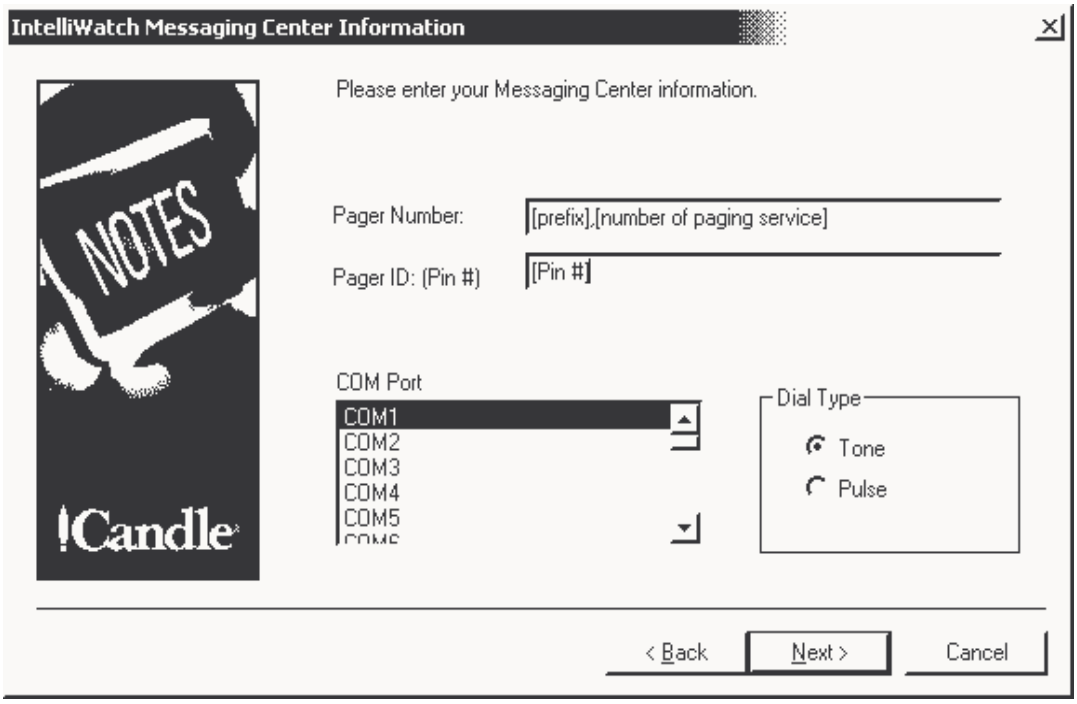

Purpose of dialog:

To allow you to configure the Paging portion of the IntelliWatch Messaging Center.

Remember ...

Don't forget to include any required prefix in the Page Number field, following by a comma.

Once Paging settings are verified, click Next to proceed.

Upgrade path:

Assuming all Pinnacle 99 components were installed on the system in question, go to *E-31*.

FIGURE E-31: Tivoli Server Location

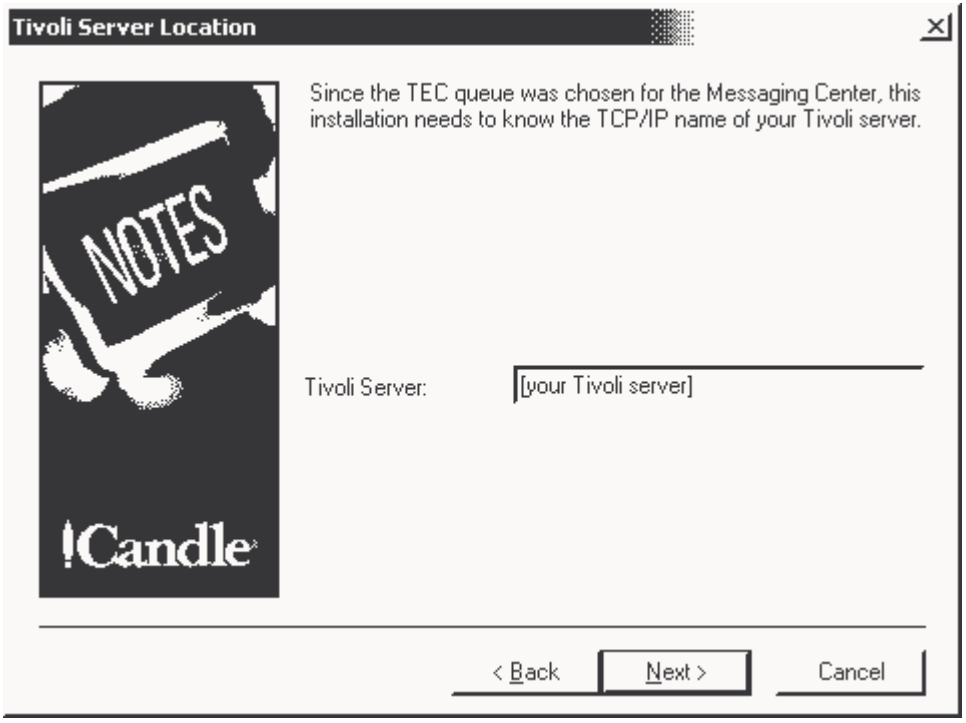

Purpose of dialog:

To allow you to enter the TCP/IP name of your Tivoli server.

This dialog is displayed only if you elected to enable the TEC Event queue at Figure E-29).

Once you have entered the name of your Tivoli server, click Next to proceed.

Upgrade path:

Assuming all Pinnacle 99 components were installed on the system in question, go to *E-35*.

FIGURE E-32: Select Advanced ServerWatch Server

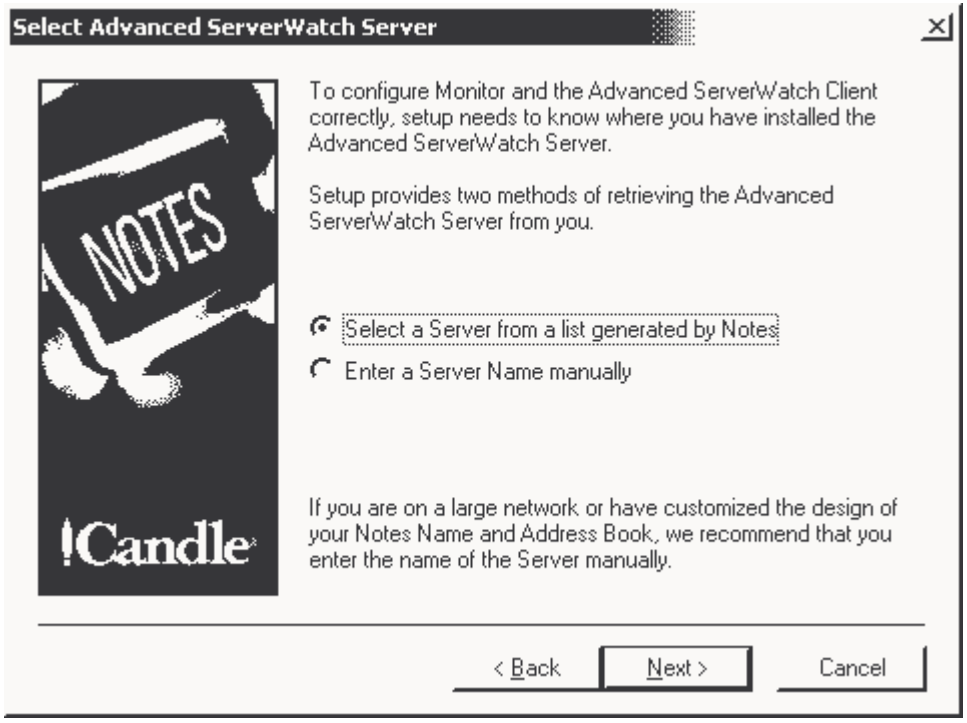

Purpose of dialog:

To allow you either to select the Advanced ServerWatch Server from a list generated by Notes, or to enter the server name manually.

Once you have made your selection, click Next to proceed.

FIGURE E-33: IntelliWatch Pinnacle: Advanced ServerWatch Hub Server

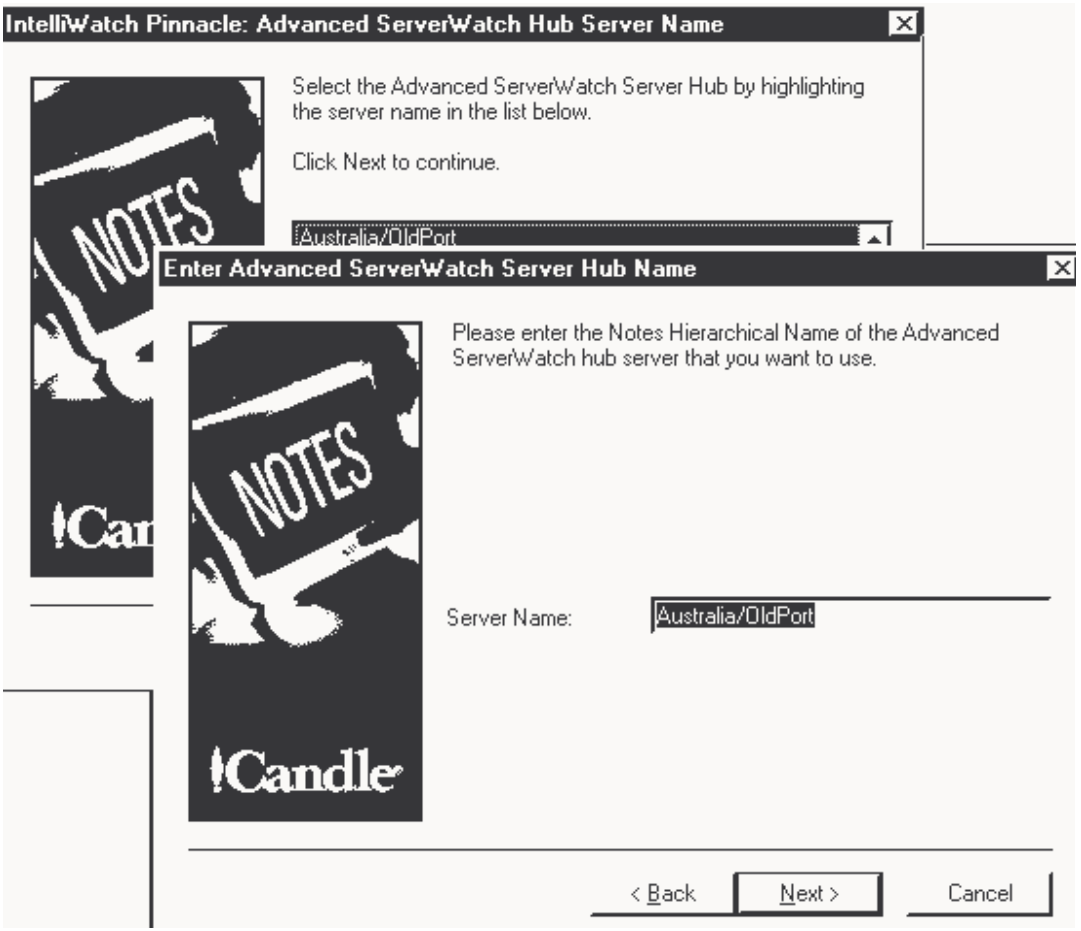

Purpose of dialog:

To allow you either to select the Advanced ServerWatch Hub Server from a list generated by Notes, or to enter the server name manually.

If you selected the first option, the dialog in the background, above, is displayed. Click on the server to select it.

If you opted to enter the server name manually, the dialog in the foreground is displayed. Enter the server name in the text box.

In either case, click Next to proceed.

FIGURE E-34: Set Program Folder

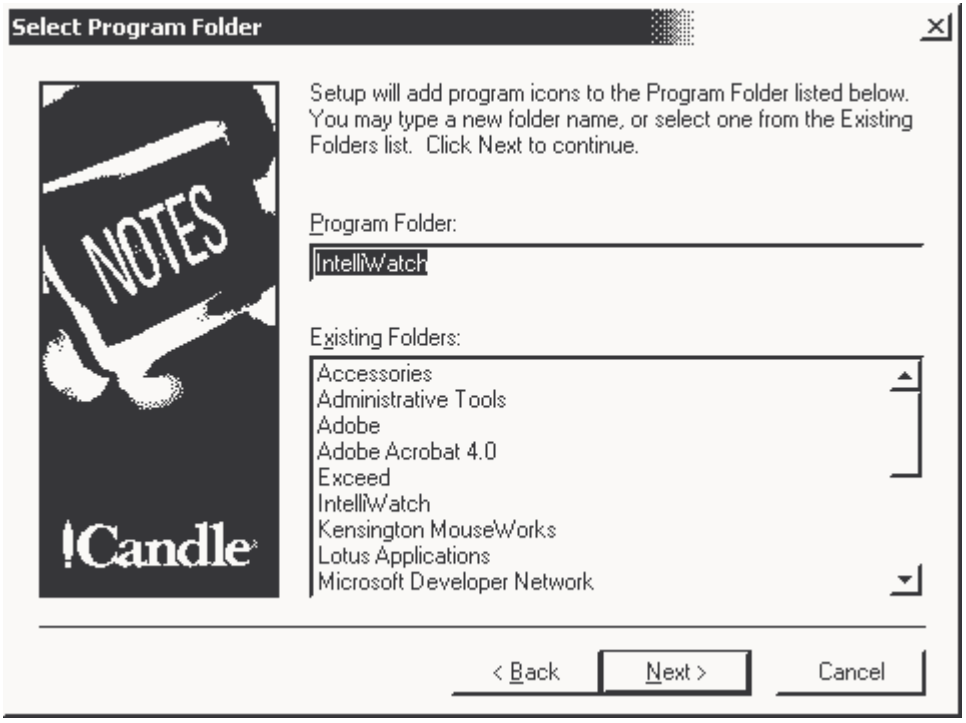

Purpose of dialog:

To allow you to customize the name of the Program Folder, or to select an existing folder. Program icons are added to the folder of your choice. (Click Next to accept the default of IntelliWatch.)

Once the correct Program Folder is displayed in the text box, click Next to proceed.

FIGURE E-35: Copying Program Files

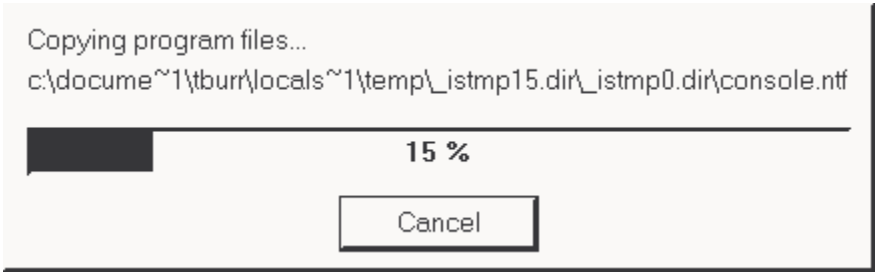

Purpose of dialog:

To allow you to follow the progress of the installation of program files.

If you want the installation to complete, do nothing. When all files have been copied, the next dialog will be displayed.

Upgrade path:

Assuming all Pinnacle 99 components were installed on the system in question, go to *E-36*.

FIGURE E-36: IntelliWatch Product Updates

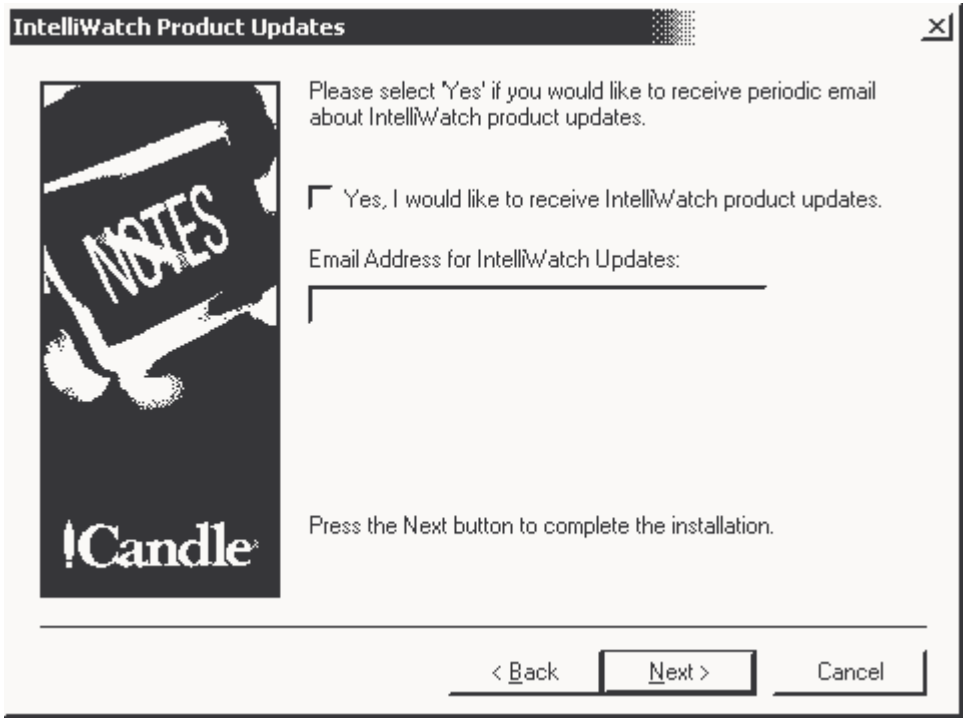

Purpose of dialog:

To allow you to request IntelliWatch product updates, as they become available.

If you want to receive product updates, select the checkbox with the mouse, then type your e-mail address into the text box.

Once you have completed your entries--or to accept the default of no e-mail updates--click Next.

Upgrade path:

Assuming all Pinnacle 99 components were installed on the system in question, go to *E-37*.

FIGURE E-37: IntelliWatch Configuration Wizard

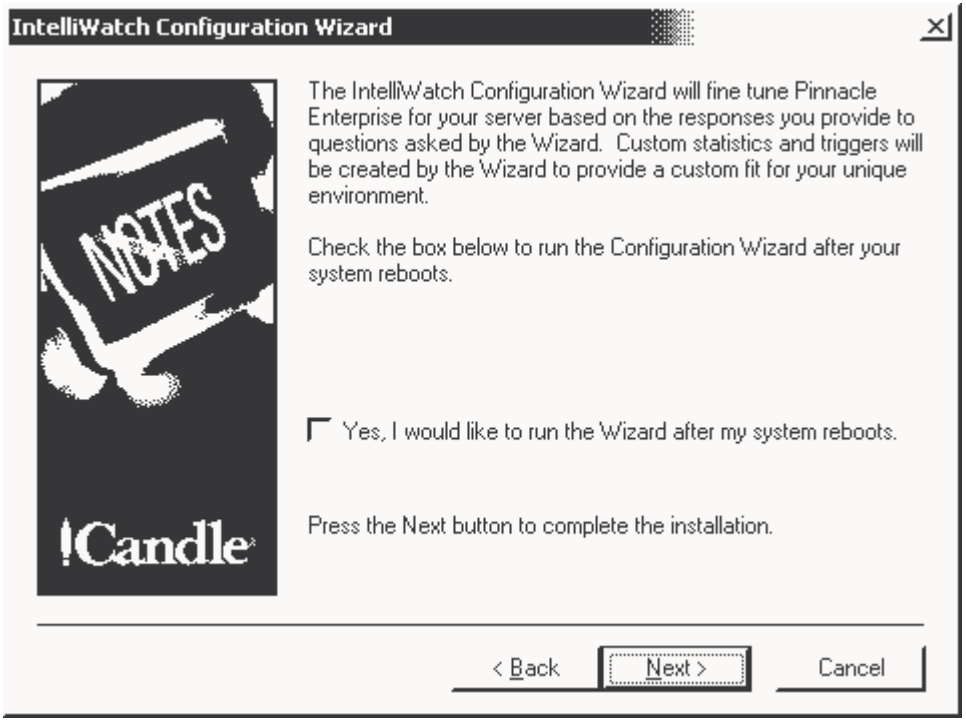

Purpose of dialog:

To allow you to launch the IntelliWatch Configuration Wizard after the system reboots.

If you want to run the Wizard, select the checkbox with the mouse.

Once you have made your selection, click Next to proceed.

Upgrade path:

Assuming all Pinnacle 99 components were installed on the system in question, go to *E-38*.

FIGURE E-38: IntelliWatch Pinnacle Setup Complete

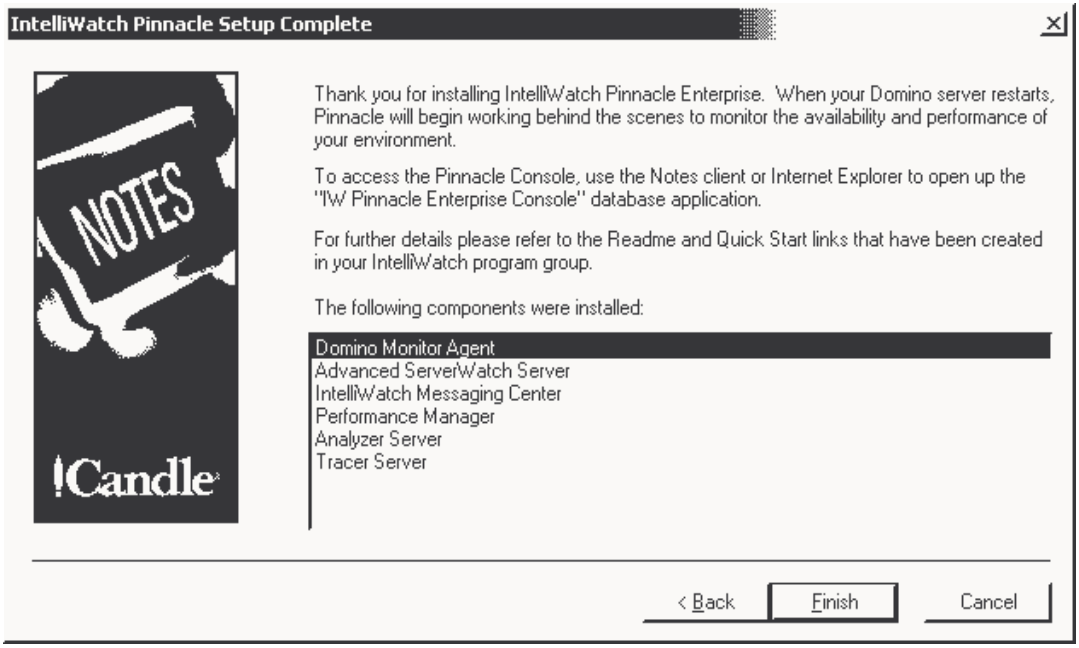

Purpose of dialog:

To display a list of the Pinnacle components that were installed by the Setup.

To complete the Setup, click Finish.

Upgrade path:

Assuming all Pinnacle 99 components were installed on the system in question, go to *E-39*.

FIGURE E-39: Restart Windows

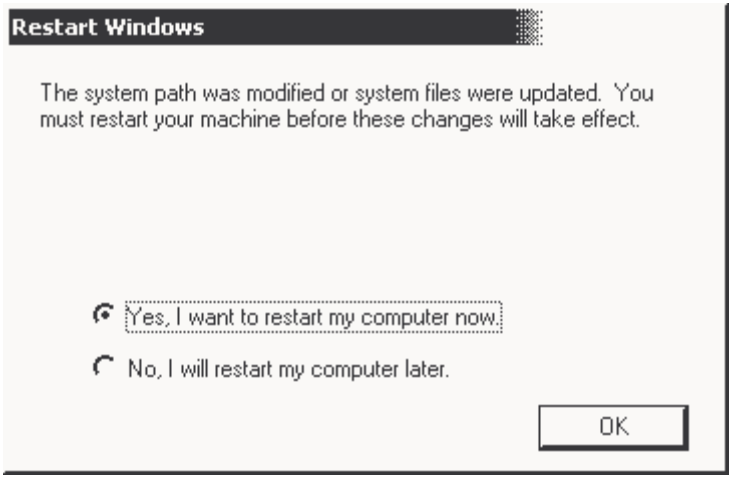

Purpose of dialog:

To allow you to select when the system reboots.

Because of changes to the system path made by the Setup, you must reboot the system before running IntelliWatch.

Once you have made your selection, click OK to complete the Setup.

FIGURE E-40: Upgrade Applications

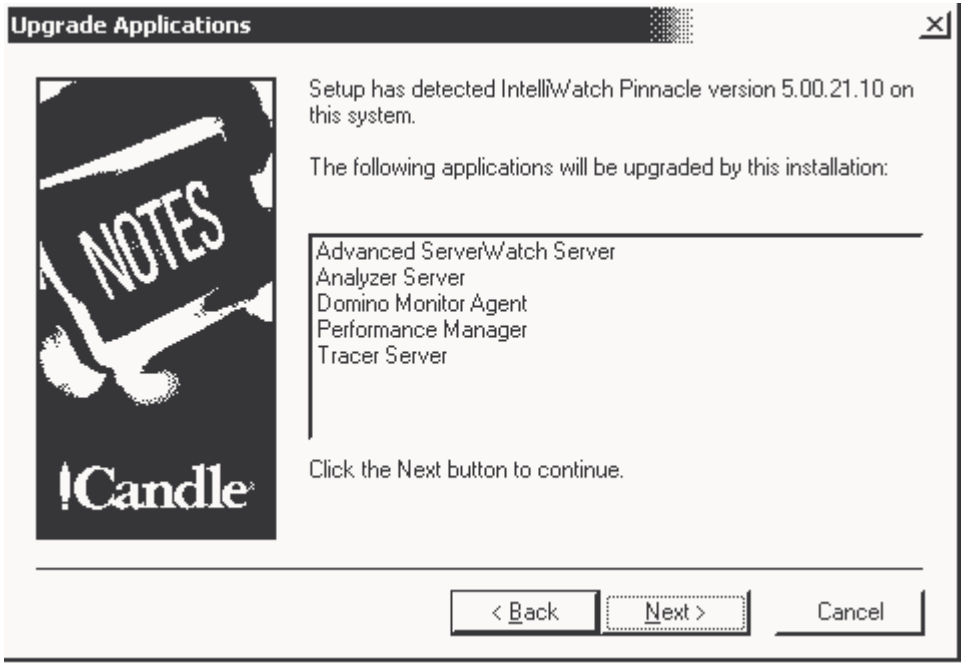

Purpose of dialog:

To display a list of the Pinnacle 99 components (version 5.00.21.10 only) that were upgraded by the Setup.

Upgrade path:

Assuming all Pinnacle 99 components were installed on the system in question, go to *E-11*.

Which Setup dialogs are displayed during an upgrade, as well as the order in which they are presented, depends on which Pinnacle 99 components were installed on the system being upgraded.

Assuming that all Pinnacle components were installed on a given system, dialogs are presented in the following order:

E-1 through E-6; E-10; E-40; E-11; E-41; E-14; E-17 and E-18; E-24 and E-25; E-28 through E-31; E-35 through E-39.

FIGURE E-41: Location of Directories

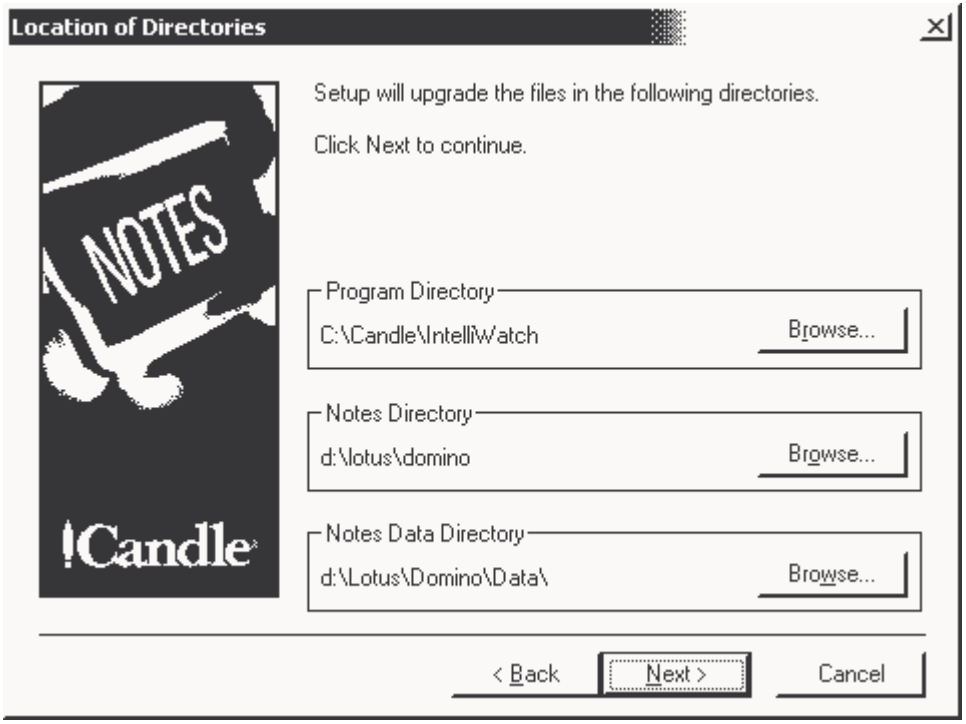

Purpose of dialog:

To inform you of the location of the Pinnacle files that will be upgraded by the Setup.

Unlike a new installation, certain choices are not available when upgrading Pinnacle components, due to decisions made during the previous installation (file locations, for instance).

You may configure products not previously on the system, however.

Upgrade path:

Assuming all Pinnacle 99 components were installed on the system in question, go to *E-14*.

Appendix

IntelliWatch Pinnacle Configuration Wizard

F

The Setup gives you the option of running the IntelliWatch Pinnacle Configuration Wizard immediately after the system reboots.

The Wizard first detects the function of the local server in your Notes environment. Based on your answers to a series of questions, it then tailors Pinnacle components to most effectively manage your Domino servers. Server-function options are:

- Application
- Hub
- Mail
- Web

Please bear in mind, however,...

that the primary purpose of the Configuration Wizard is as a tool to be used by on-site SEs (to automate product demonstrations).

Therefore, be aware that running the Wizard as part of a normal installation creates a significant likelihood that you will want (or need) to later disable (or modify) much of what the Wizard generates.

FIGURE F-1: IntelliWatch Pinnacle Configuration Wizard

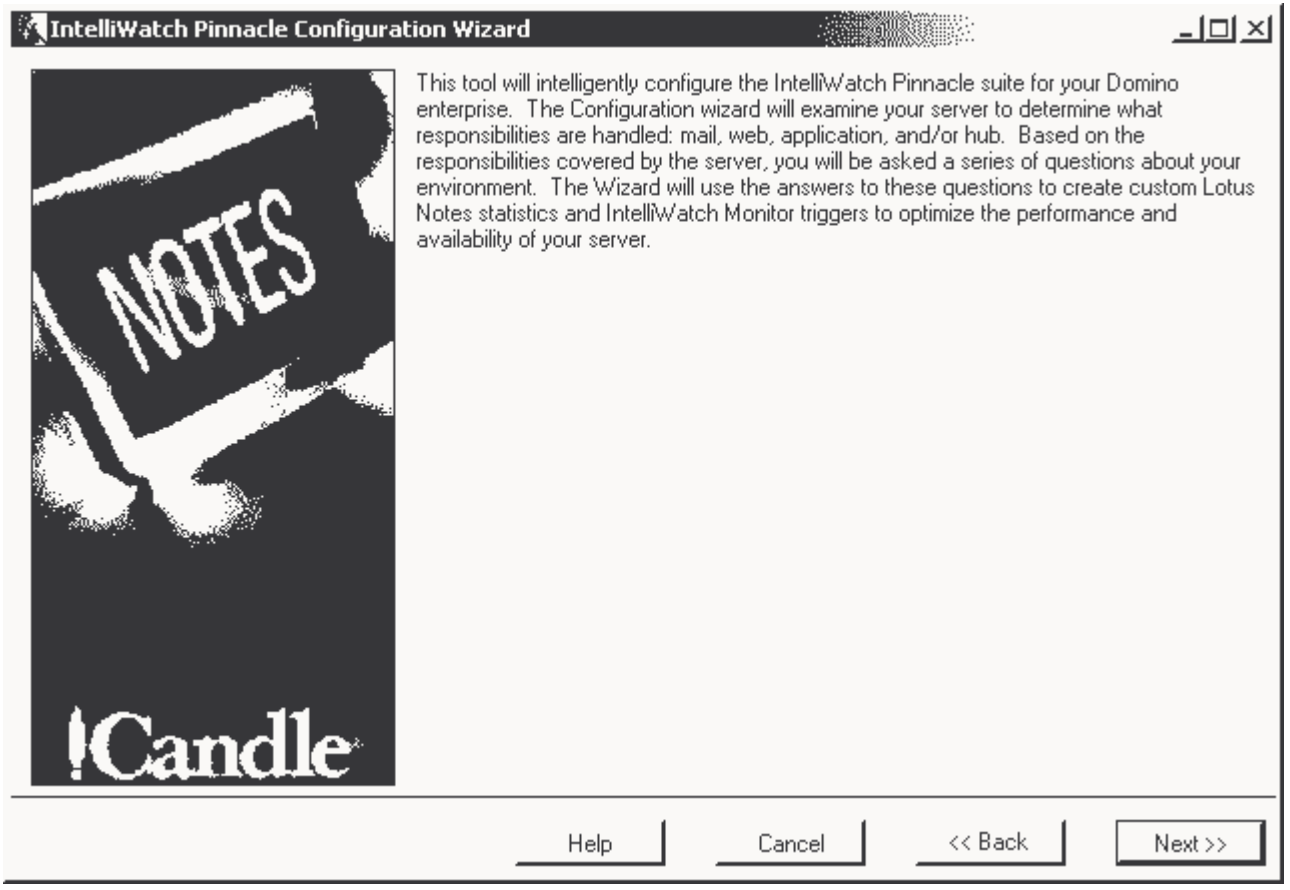

FIGURE F-2: Configuration Wizard: Server Responsibilities

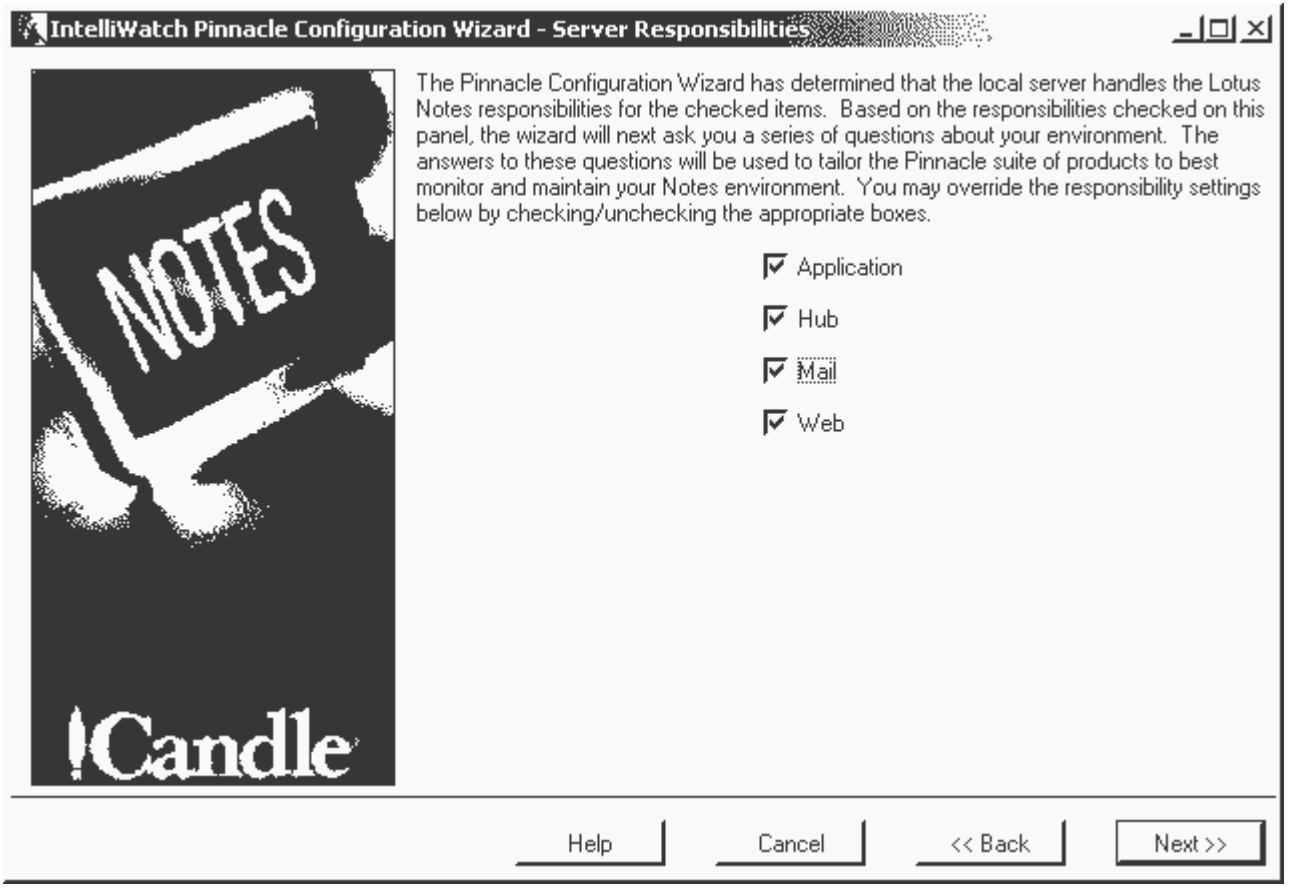

FIGURE F-3: Configuration Wizard: Trigger Notification

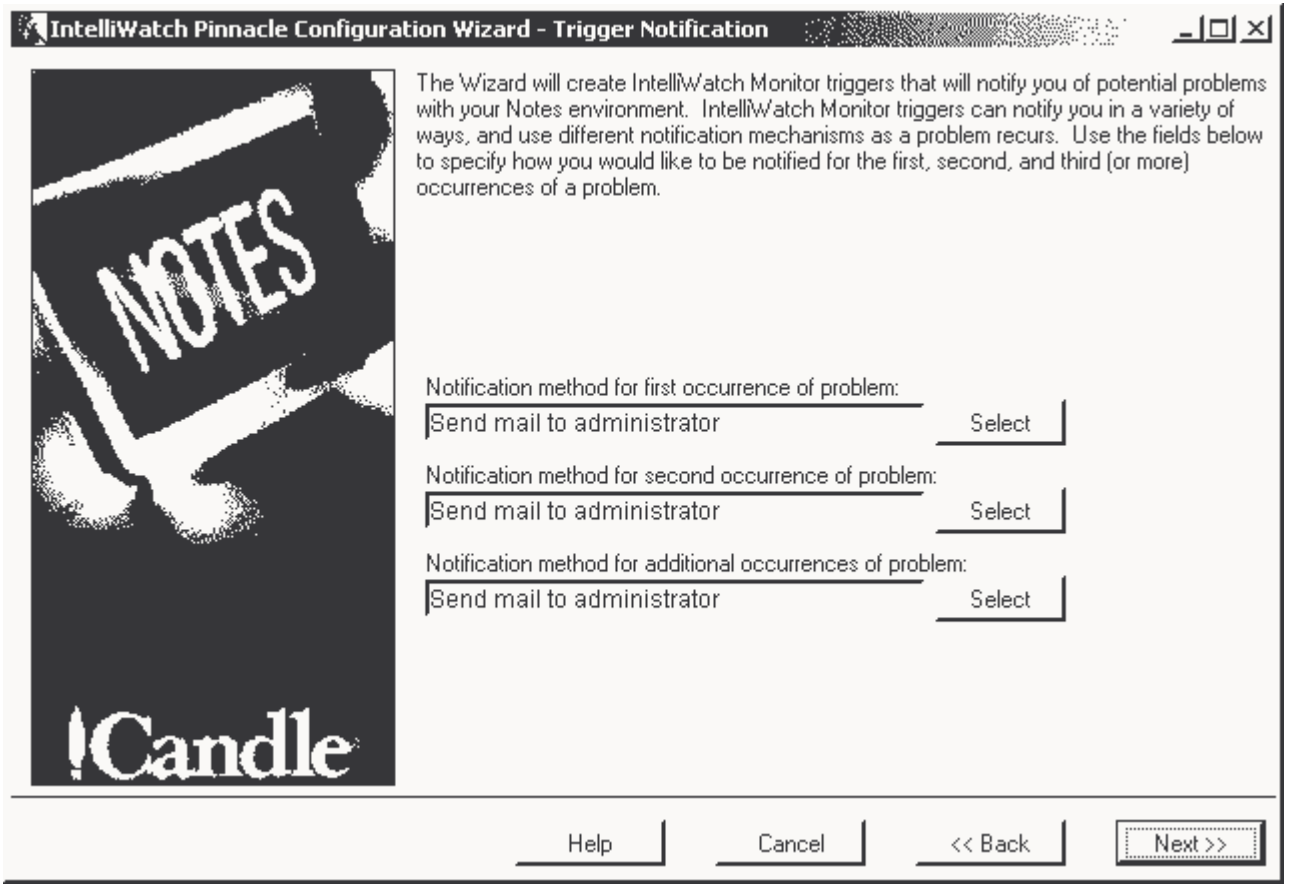

FIGURE F-4: Server: Database Performance and Availability

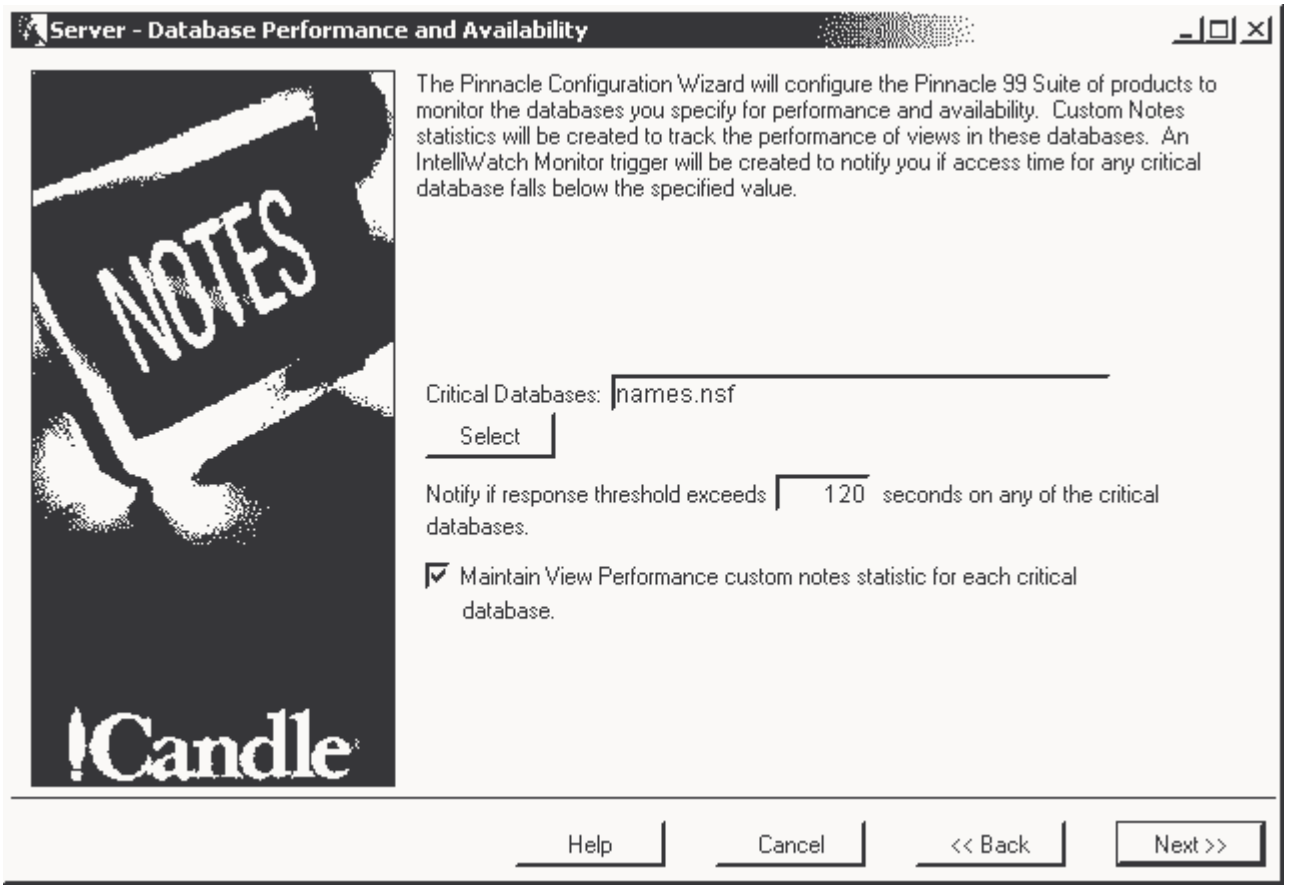

FIGURE F-5: Server: Database Activity and Security

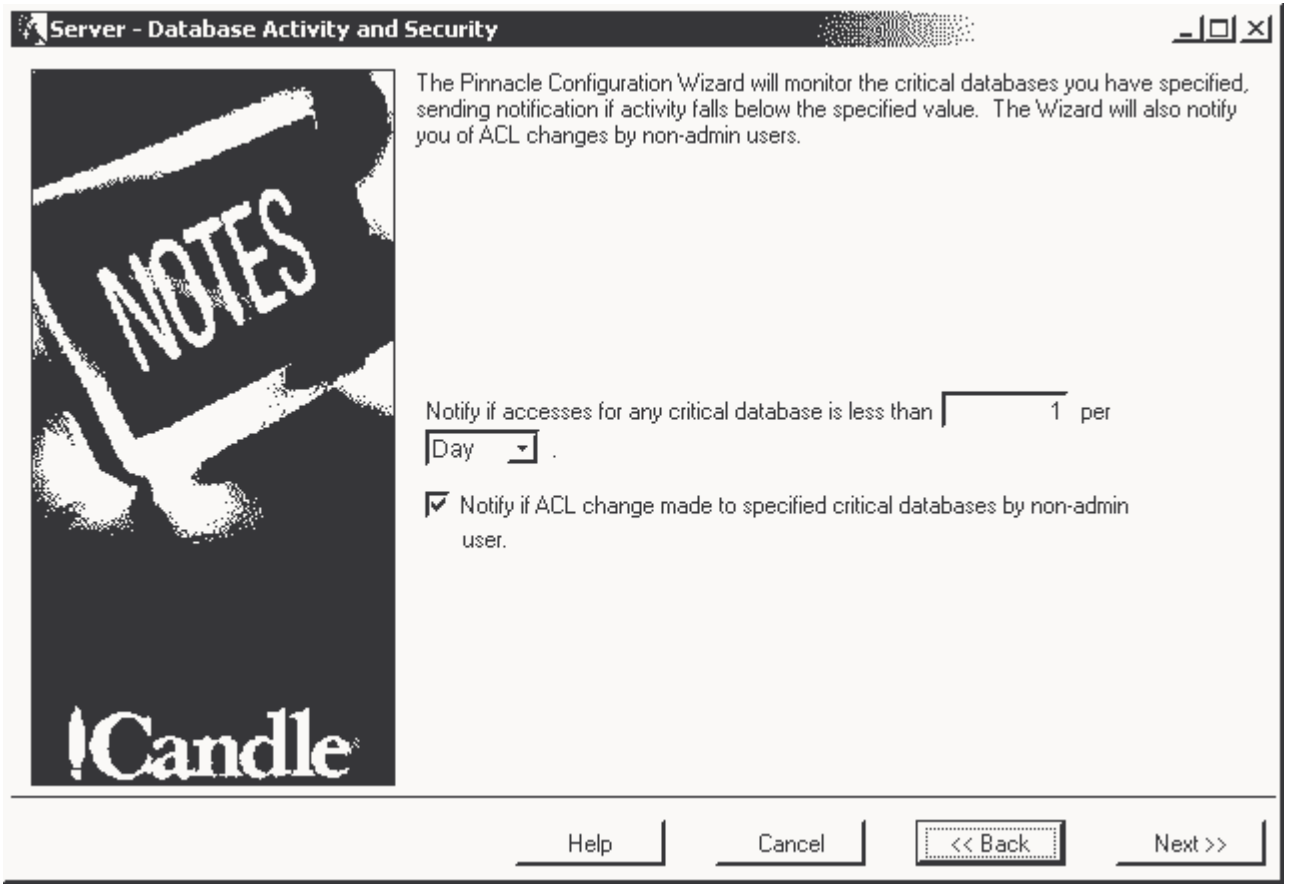

FIGURE F-6: Mail Server: Availability and Performance

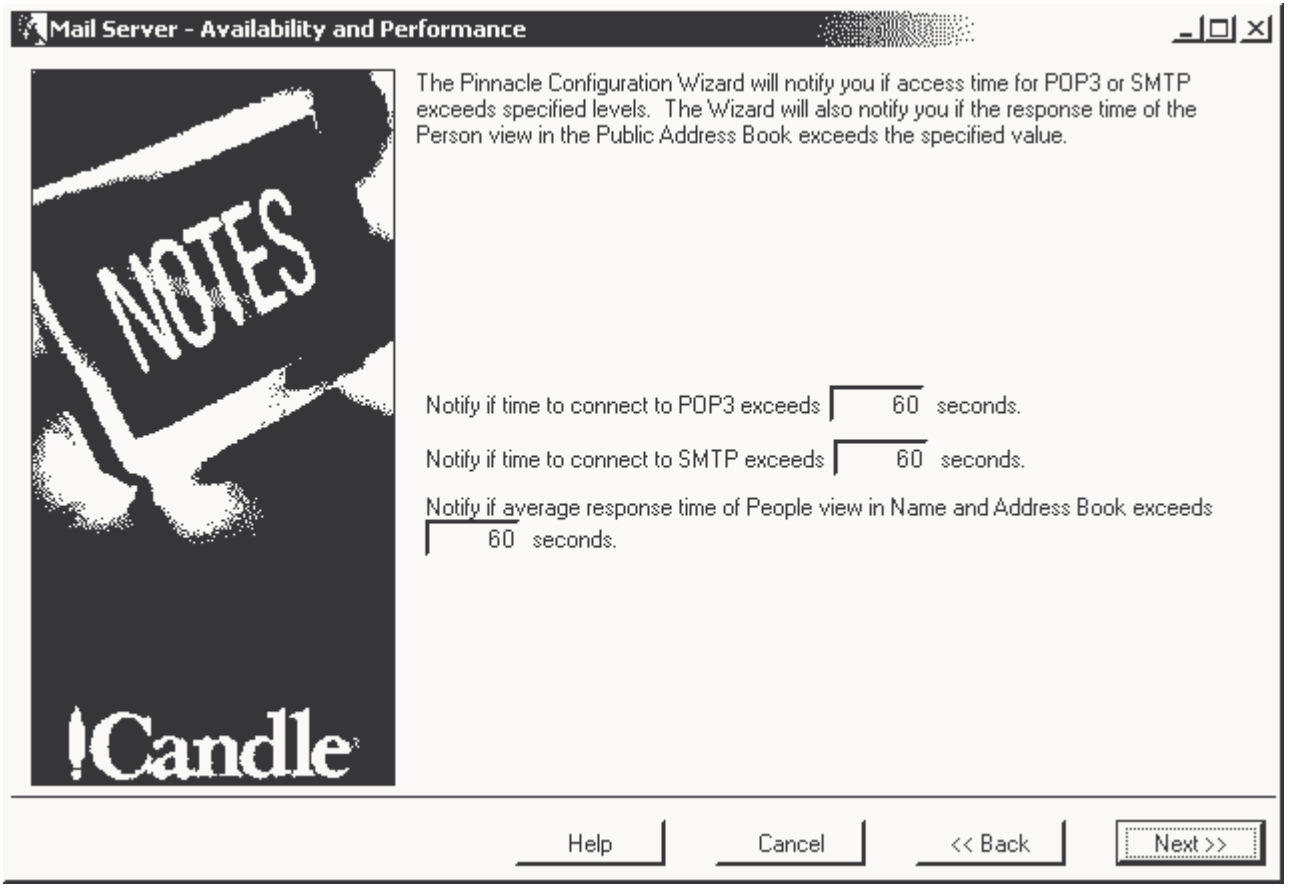

FIGURE F-7: Mail Server: Availability and Maintenance

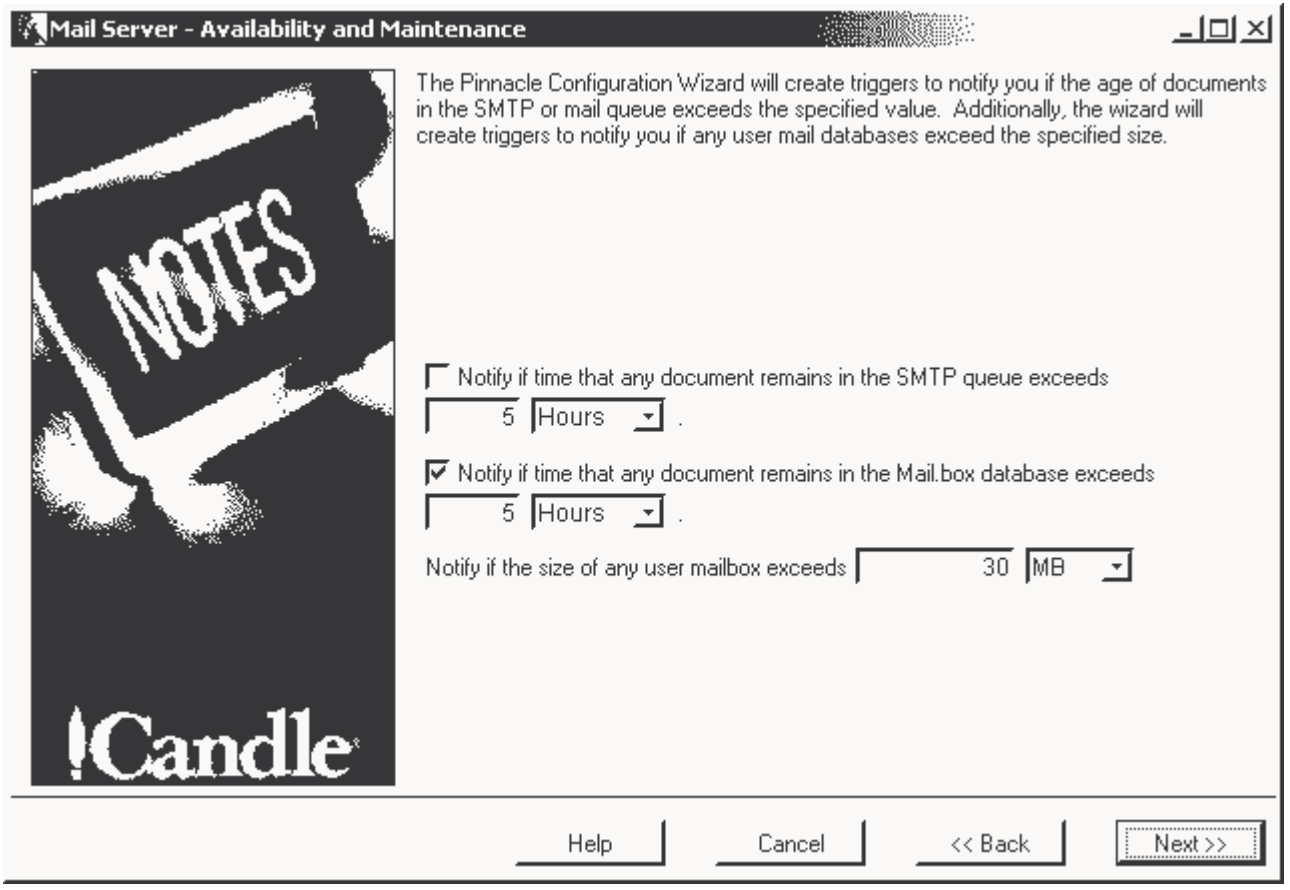

FIGURE F-8: Mail Server: Availability

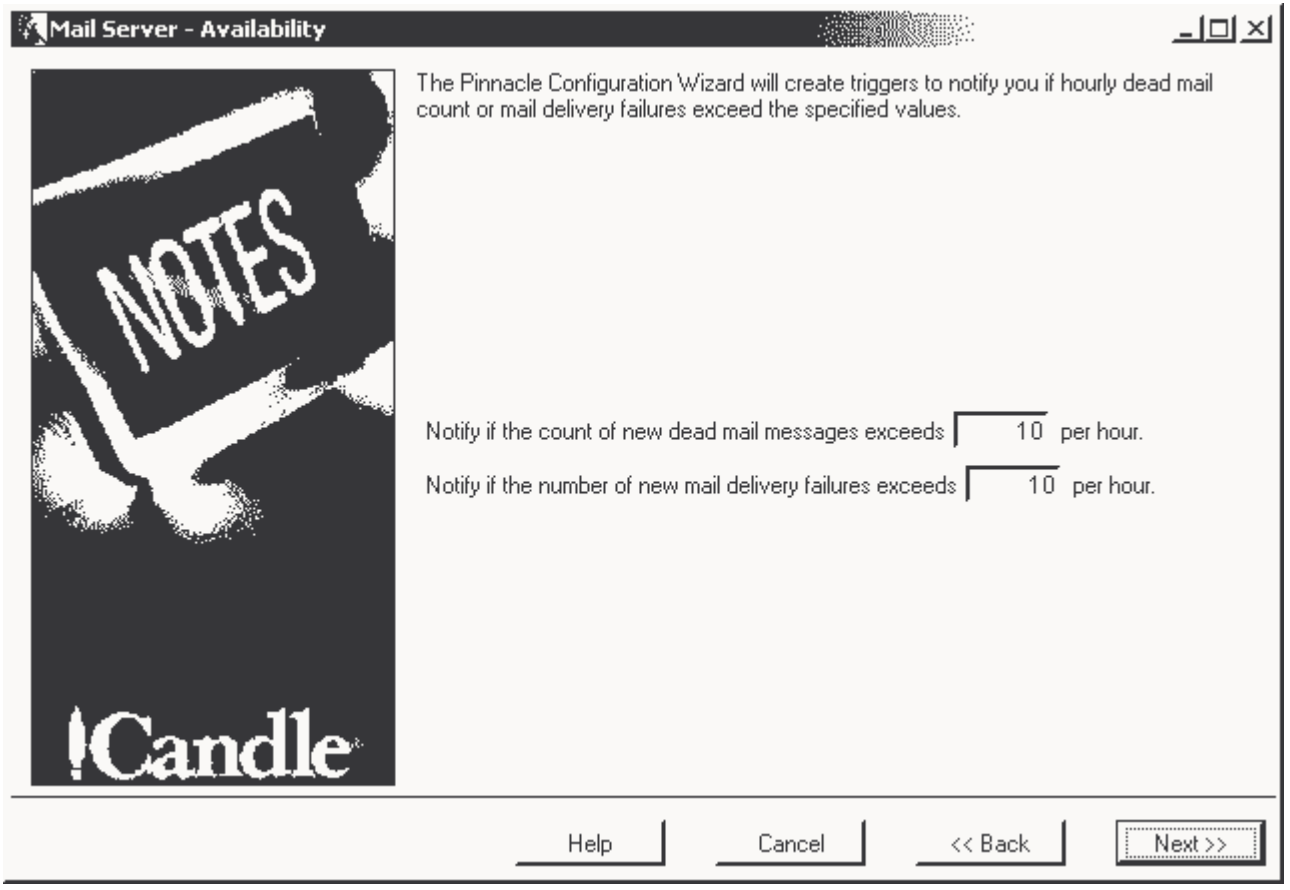

FIGURE F-9: Web Server: Availability and Performance

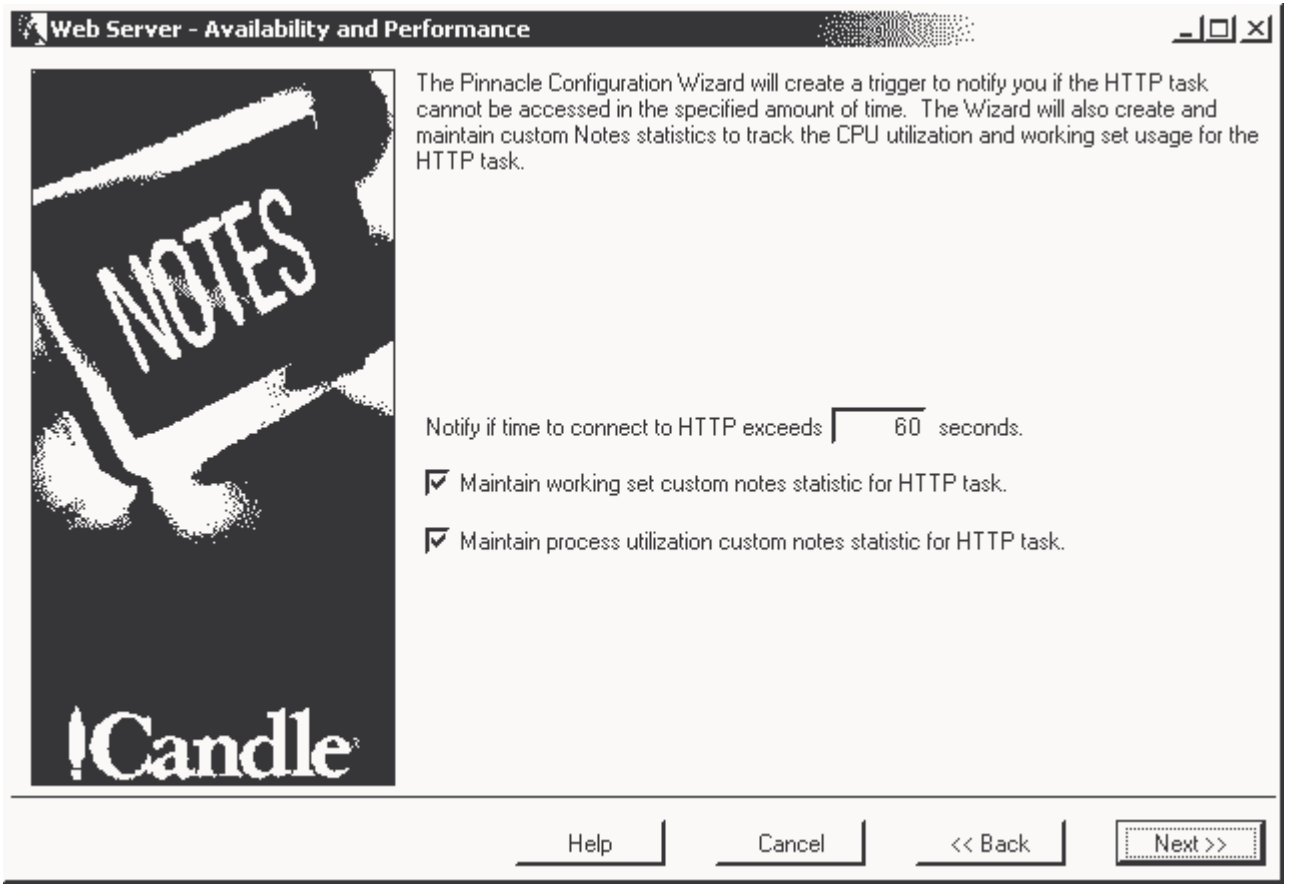

FIGURE F-10: Hub Server: Database Replication

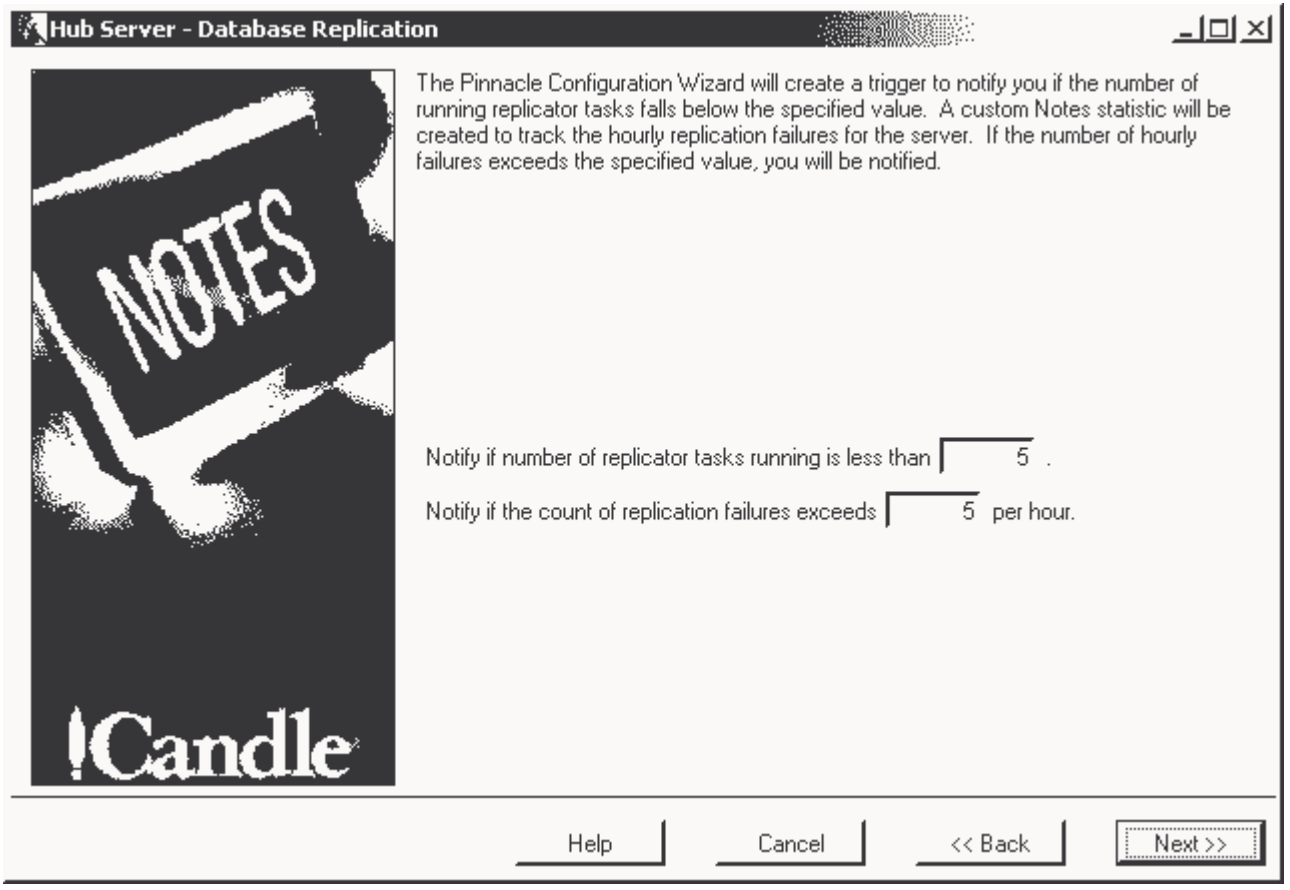

FIGURE F-11: Hub Server: Database Replication (continued)

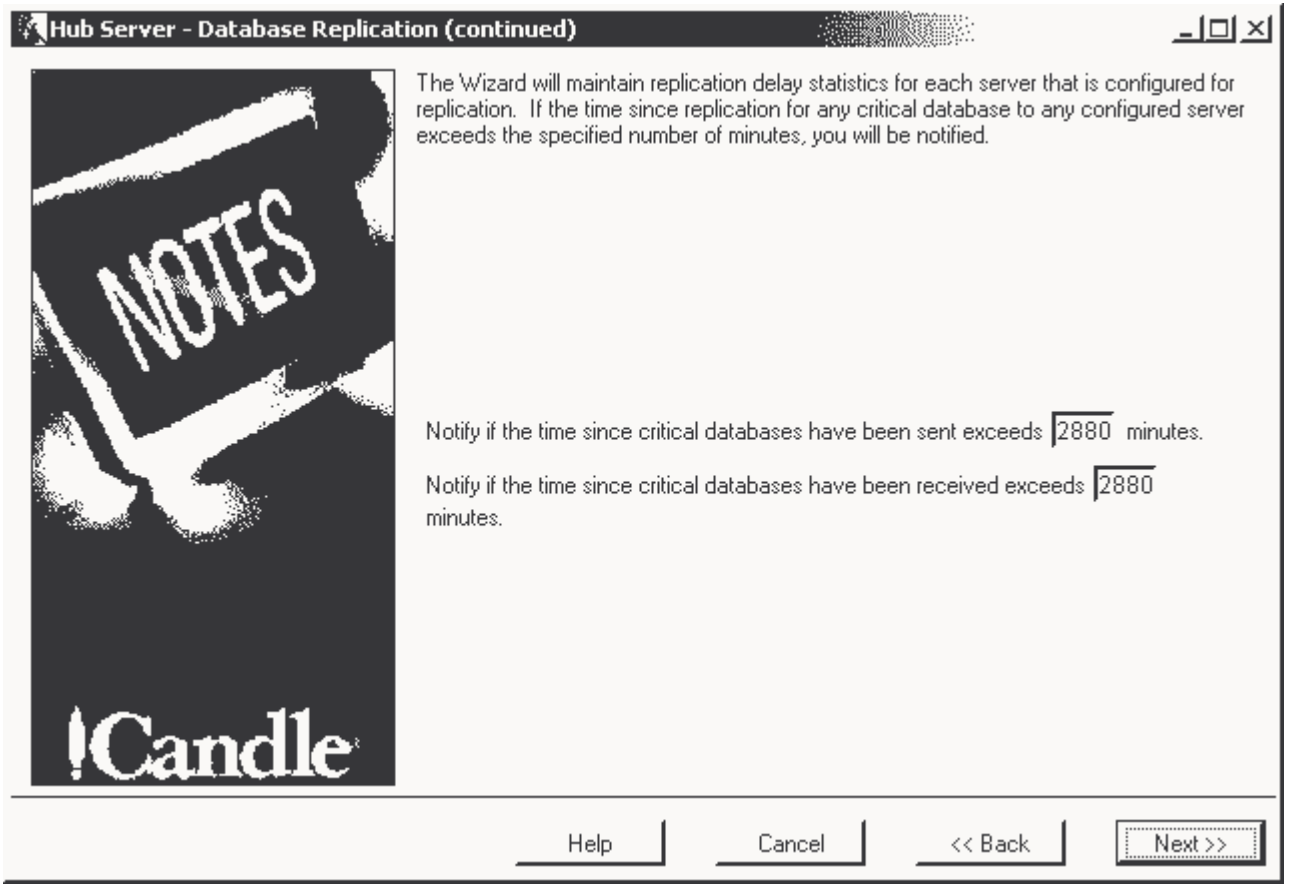

FIGURE F-12: Application Server: Workflow

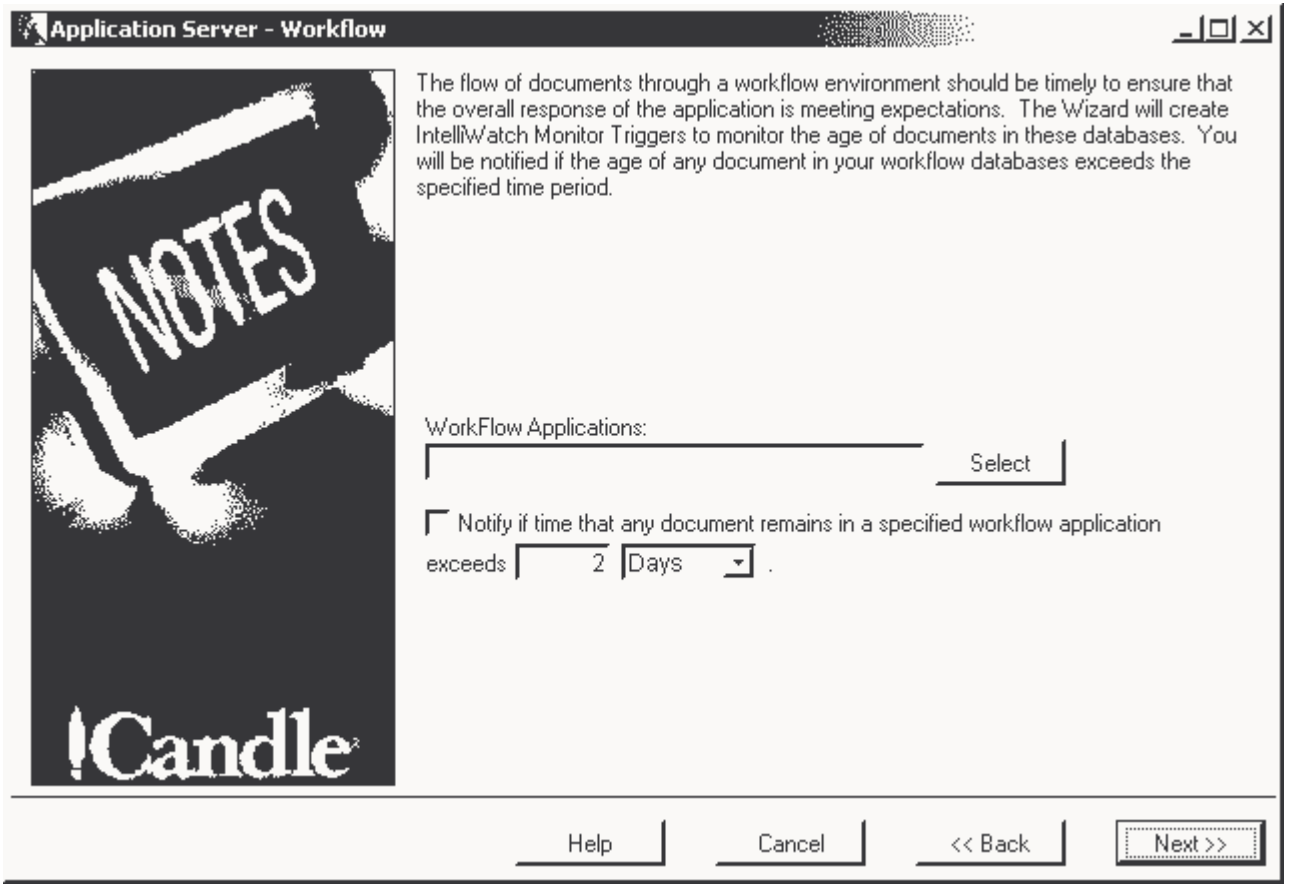

FIGURE F-13: Configuration Wizard: Server Groups

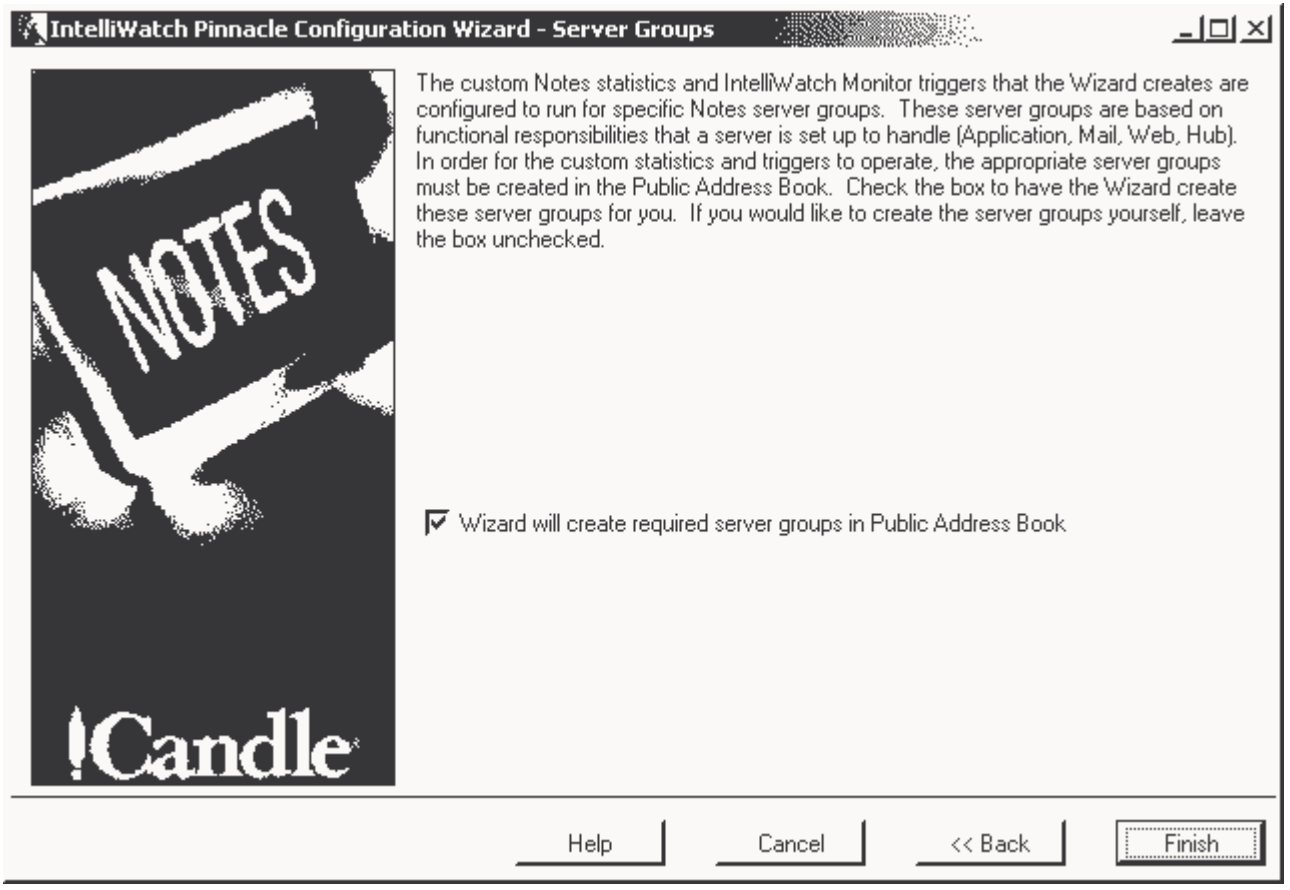

FIGURE F-14: Create Triggers and Statistics

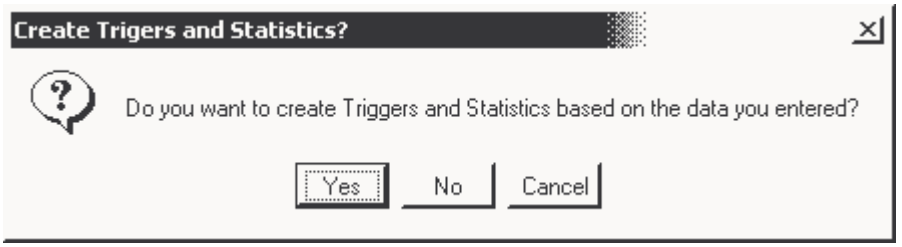

FIGURE F-15: Configuring System

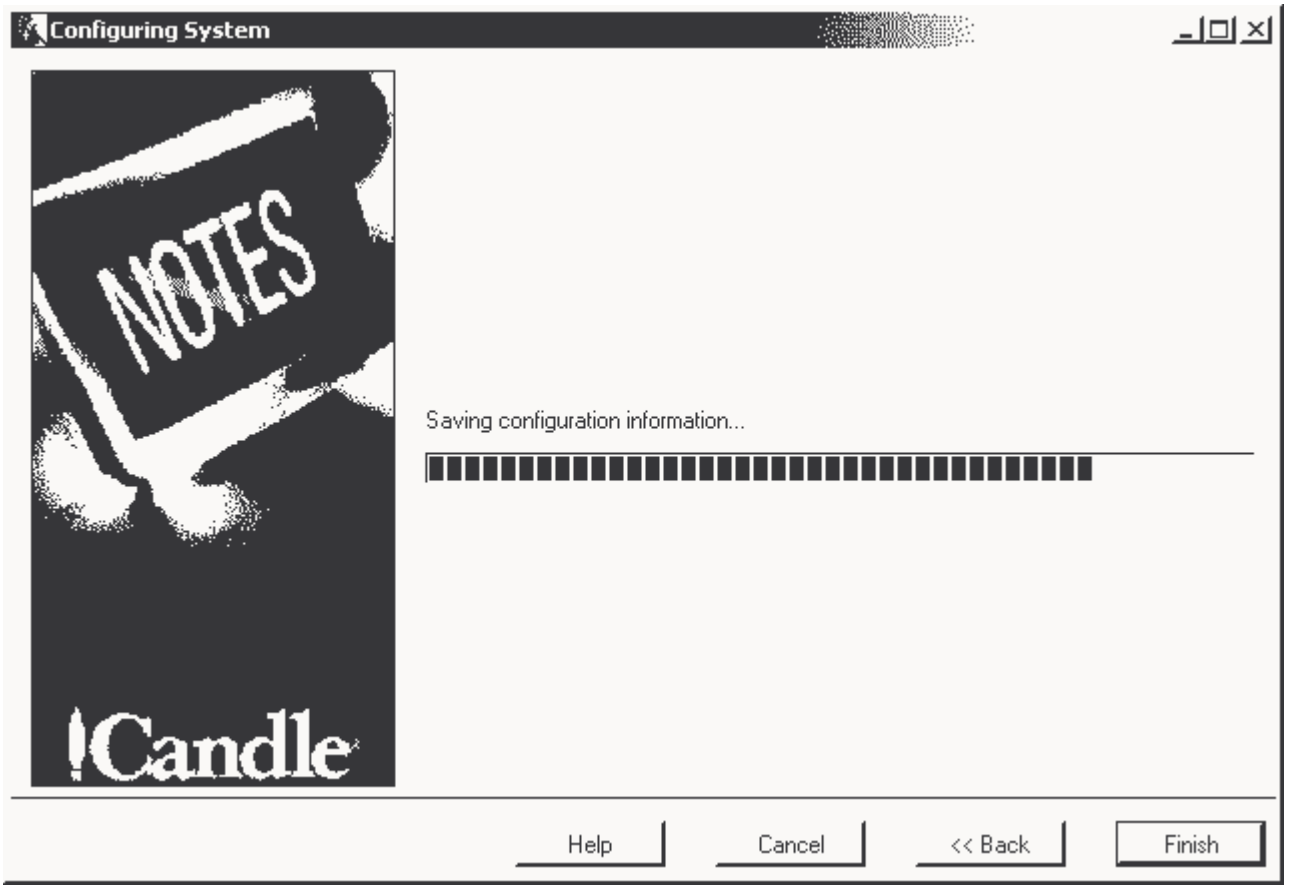

FIGURE F-16: Configuration Wizard: Completion Screen

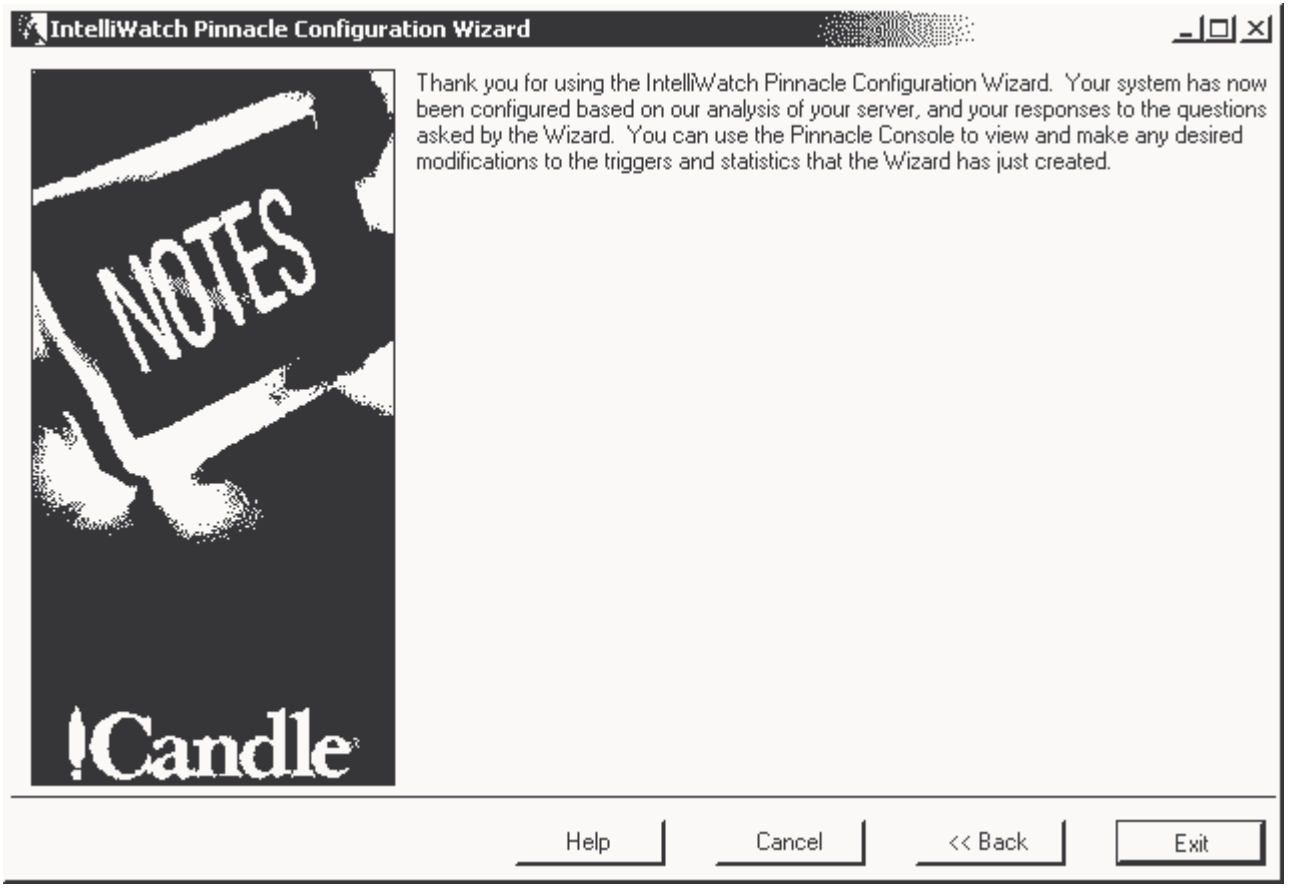

Appendix

Data Returned by PM Statistic Types

G

As discussed in *"Chapter 5, Performance Manager"* (under *5.1.0.0*), each PM statistic type returns very specific data.

As an aid to understanding what is returned by each type, the following table is provided.Data returned by PM statistic types.

Table G-1. Data Returned by PM Statistic Types

PM Statistic Type	Returned (in addition to Statistic Name)
Average	No additions to statistic name.
Delta	No additions to statistic name (including time interval)
Difference	No additions to statistic name.
File System Information	UNIX only: [mount point].Free [mount point].Size
Mail Domain	incoming.[domain name].Count incoming.[domain name].Size outgoing.[domain name].Count outgoing.[domain name].Size
Mail Incoming Delivery	[server name].CountOver and/or [server name].Max [server name].Avg [server name].Min
Mail Outgoing Attachment Types	[extension].Count [extension].Size
Mail Outgoing Server Attachment Percentage	No additions to statistic name (for overall percentage) Overall percentage broken down by: [destination server name]
Mail Outgoing Server Volume	[server name].Count [server name].Volume (in bytes)

Table G-1. Data Returned by PM Statistic Types

PM Statistic Type	Returned (in addition to Statistic Name)
Mail Size	**USER LEVEL:** Sent.[User, i.e. sender].Max Sent.[User, i.e. sender].Avg Sent.[User, i.e. sender].Count Received.[User, i.e. recipient].Max Received.[User, i.e. recipient].Avg Received.[User, i.e. recipient].Count **SERVER LEVEL:** Outgoing.[server name].Max Outgoing.[server name].Avg Outgoing.[server name].Count **MAXIMUM/AVERAGE SUMMARY:** Sent.Max Sent.Avg Sent.Count Recv.Max Recv.Avg Recv.Count
NT Performance Counter	No additions to statistic name.
Replication	Pulled.KBReceived.[server name] Pulled.KBSent.[server name] Pulled.Documents.Transferred.[server name] Pushed.KBReceived.[server name] Pushed.KBSent.[server name] Pushed.Documents.Transferred.[server name] Time.[server name] TotalMissed.[server name] TotalSuccesses.[server name]
Replication Delay	[server name].MinutesSince.Recv [server name].MinutesSince.Send
Server Event Count	No additions to statistic name.
Summation	No additions to statistic name.

Table G-1. Data Returned by PM Statistic Types

PM Statistic Type	Returned (in addition to Statistic Name)
View Performance	**SERVER AND DATABASE SUMMARIES:** Avg Count Max PercentUnder[SLA threshold from Type Information tab] **BY VIEW:** [View].Avg [View].Count [View].Max [View].PercentUnder[SLA threshold from Type Information tab] **BY VIEW (WHEN VIEW IN A FOLDER):** [Folder].\[View].Avg [Folder].\[View].Count [Folder].\[View].Max [Folder].\[View].PercentUnder[SLA threshold from Type Information tab]

Appendix

Ports Used by IntelliWatch

H

Port basics (non-partitioned servers):

IntelliWatch port usage on non-partitioned Domino servers is relatively straightforward. Although users are free to customize certain port settings, most users simply accept the defaults, which are listed in the following table. (For information on what may and may not be customized, see the section Port customization, below.):

Table H-1. Ports used by IntelliWatch

Port usage	*Port #*
IWASW Server Task	**22576**
IWINFO Server Task	**23576**
Message Center Gateway (paging requests)	**21370**
Message Center Gateway (all other requests)	**21371**
Portmapper (running as NT Service on Windows, and as a process on UNIX)	**24576**
Tracer Server Task (This task uses the first available port from the following list.)	**15130** **16150** **17380** **18930** **19390** **20580** **21630** **22510** **23930** **24780**

Message Center Gateway

The Message Center Gateway listens on ports 21370 (paging) and 21371 (other notifications).

Any Pinnacle component or utility that sends something via the Message Center Gateway does so via these ports (iwpage, iwevent, iwasw, iwagent paging, etc.).

ASW and the Notes ping:

ASW checks server connectivity by means of a Notes ping. ASW does not necessarily use port 1352 for this purpose, but instead issues a command using the Notes API--which causes Notes to talk to the other server by means of some underlying Network protocol.

This protocol might be TCP/IP (which by default uses port 1352), or some other network protocol, depending on how Notes is configured in a given environment. (In the default situation, port 1352 would need to be open to the ASW Hub system on all servers that the Hub will be monitoring.)

In short, iwasw's pinging of the server depends on underlying Notes communication.

Port basics (partitioned servers):

Portmapper:

Portmapper exists for the purpose of coordinating communication between the Pinnacle Console and IntelliWatch tasks running on partitioned servers.

Running as an NT Service on Windows or as a process on UNIX, Portmapper listens on port 24576. When the Pinnacle Console connects to a server via this port, communication proceeds as follows:

1. The Pinnacle Console connects to port 24576 on the desired host (server short name or the hostname specified in the UI).
2. Portmapper accepts the connection(s) and reads from the socket.
3. The Console sends the name of the server it wants to talk to.
4. Portmapper reads the server name from the socket and looks up its partition number in the NT registry or in the iwglobal.ini file on UNIX:

```
NT registry: (HKEY_LOCAL_MACHINE\SOFTWARE\Candle\IntelliWatch\[Server Name]\Partition=[#])
iwglobal.ini: [Server Name]
                Partition=[#]
```

5. Portmapper sends the partition number down the socket.
6. The Pinnacle Console reads the number and disconnects.
7. Portmapper disconnects as well, performs some cleanup and then goes back to waiting for someone to connect.
8. The Console adds the number it is given to the base port for the task in question, then establishes a connection with the task, whereupon it carries out the required action(s).

IWINFO and ASW

On partitioned servers, the port numbers in the above list are used by default, and are then incremented by the partition number.

IWINFO, running on the second partition of a system, for example, would listen on 23576 + 2, or port 23578.

Note: IWINFO doesn't initiate communication with any other process (except the Message Center Gateway, when sending a page).

IWTRACE

Tracer listens on the first port it finds to be available from the above list, then offsets it by the partition number in the same fashion as IWINFO and IWASW.

The part of Tracer installed on the Notes client listens on the same set of ports, and a partition number of 0 is assumed. The Enterprise console must be able to connect to the port Tracer is listening on, on both the server and the client. After the initial connection is made, the OS opens a new port (not currently in use), and the Pinnacle Console reconnects to that port.

Port customization:

ALL ports above can be set to just about anything above 1024 via either the NT registry (on Windows systems) or in the *.ini files (on UNIX). There are two exceptions:

- Tracer ports
- The TCP/IP Notes port (which is governed by a setting in the notes.ini file)

Server NIC key:

In tandem with the above port settings, all IntelliWatch Pinnacle components use the following Server NIC setting to determine which network interface to listen on:

NT registry: HKEY_LOCAL_MACHINE\SOFTWARE\Candle\IntelliWatch\Network\Server NIC=[ip address]
Iwpinn.ini: [Network]
Server NIC=[ip address]

If this setting is not present (please note that it must be created manually), Pinnacle attempts to listen on all interfaces. In some environments, however, the server components are only allowed to listen on one interface. On systems with multiple NICs, this can lead to connectivity issues. Adding this key resolves most, if not all, such issues.

Appendix

Support for problem solving

If you have a problem with your IBM software, you want to resolve it quickly. This section describes the following options for obtaining support for IBM software products:

Using IBM Support Assistant

The IBM Support Assistant is a free, stand-alone application that you can install on any workstation. You can then enhance the application by installilng product-specific plug-in modules for the IBM products you use.

The IBM Support Assistant saves you the time it takes to search the product, support, and educational resources. The IBM Support Assistant helps you gather support information when you need to open a problem management record (PMR), which you can then use to track the problem.

The product-specific plug-in modules provide you with the following resources:

- Support links
- Education links
- Ability to submit problem management reports

For more information, and to download the IBM Support Assistant, see http://www.ibm.com/software/support/isa. After you download and install the IBM Support Assistant, follow these steps to install the plug-in for your Tivoli product:

1. Start the IBM Support Assistant application.
2. Select **Updater** on the Welcome page.
3. Select **New Properties and Tools** or select the **New Plug-ins** tab (depending on the version of IBM Support Assistant installed).
4. Under **Tivoli**, select your product, and then click **Install**. Be sure to read the license and description. If your product is not included on the list under Tivoli, no plug-in is available yet for the product.
5. Read the license and description, and click **I agree.**
6. Restart the IBM Support Assistant.

Obtaining fixes

A product fix might be available to resolve your problem. To determine which fixes are available for your Tivoli software product, follow these steps:

1. Go to the IBM Software Support Web site at http://www.ibm.com/software/support.
2. Under **Select a brand and/or product**, select **Tivoli**.
3. Select your product and click **Go**.
4. Under **Download**, click the name of a fix to read its description and, optionally, to download it. If there is no **Download** heading for your product, supply a search term, error code, or APAR number in the field provided under **Search Support (this product)**, and click **Search**.

For more information about the types of fixes that are available, see the *IBM Software Support Handbook* at http://techsupport.services.ibm.com/guides/handbook.html.

Receiving weekly support updates

To receive weekly e-mail notifications about fixes and other software support news, follow these steps:

1. Go to the IBM Software Support Web site at http://www.ibm.com/software/support.
2. Click My support in the far upper-right corner of the page under **Personalized support**.
3. If you have already registered for **My support**, sign in and skip to the next step. If you have not registered, click register now. Complete the registration form using your e-mail address as your IBM ID and click **Submit**.
4. The **Edit profile** tab displays.
5. In the first list under **Products**, select **Software**. In the second list, select a product category (for example, **System and Asset Management**). In the third list, select a product sub-category (for example, **Application Performance & Availability** or **Systems Performance**). A list of applicable products displays.
6. Select the products for which you want to receive updates. For example, IBM Tivoli IntelliWatch.
7. Click **Add products**.
8. After selecting all products that are of interest to you, click **Subscribe to email** on the **Edit profile** tab.
9. In the **Documents** list, select **Software**.
10. Select **Please send these documents by weekly email**.
11. Update your e-mail address as needed.
12. Select the types of documents you want to receive.
13. Click **Update**.

If you experience problems with the **My support** feature, you can obtain help in one of the following ways:

Online Send an email message to erchelp@ca.ibm.com, describing your problem.

By phone Call 1-800-IBM-4-You (1-800-426-4968)

Contacting IBM Software Support

IBM Software Support provides assistance with product defects. The easiest way to botain that assistance is to open a PMR or ETR directly from the IBM Support Assistant (see).

Before contacting IBM Software Support , your company must have an active IBM software maintenance contract, and you must be authorized to submit problems to IBM. The type of software maintenace contract that you need depends on the type of product you have:

- For IBM distributed software products (including, but not limited to, Tivoli, Lotus, and Rational® products, as well as DB2 and WebSphere® products that run on Windows or UNIX operating systems), enroll in Passport Advantage in one of the following ways:
  - **Online** Go to the Passport Advantage Web site at http://www-306.ibm.com/software/howtobuy/passportadvantage/pao_customers.htm
  - **By phone** For the phone number to call in your country, go to the IBM Software Support Web site at http://techsupport.services.ibm.com/guides/contacts.html and click the name of our geographic region.
- For customers with Subscription and Support (S & S) contracts, go to the Software Service Request Web site at https://techsupport.services.ibm.com/ssr/login.
- For customers with IBMLink™ CATIA, Linux, OS/390, iSeries®, pSeries®, zSeries, and other support agreements, go to the IBM Support Line Web site at http://www.ibm.com/services/us/index.wss/so/its/a1000030/dt006.
- For IBM eServer™ software products (including but not limited to, DB2 and WebSphere products that run in zSeries, pSeries, and iSeries environments), you can purchase a software maintenance agreement by working directly with an IBM sales representative or an IBM Business Partner. For more information about support for eServer software products, go to the IBM Technical Support Advantage Web site at http://www.ibm.com/servers/eserver/techsupport.html.

If you are not sure what type of software maintenance contract you need, call 1-800-IBMSERV (1-800-426-7378) in the United States. From other countries, go to the contacts page of the IBM Software Support Handbook on the Web at http://techsupport.services.ibm.com/guides/contacts.html and click the name of your geographic region for phone numbers of people who provide support for your location.

To contact IBM Software Support, follow these steps:

1. Determining the business impact
2. Describing problems and gathering information
3. Submitting problems

Determining the business impact

When you report a problem to IBM, you are asked to supply a severity level. Use the following criteria to understand and assess the business impact of the problem you are reporting.

Severity 1 The problem has a critical business impact. You are unable to use the program, resulting in a critical impact on operations. This condition requires an immediate solution.

Severity 2 - The problem has a significant business impact. The program is usable, but it is severely limited.

Severity 3 The problem has some business impact. The program is usable, but less significant features (not critical to operations) are unavailable.

Severity 4 The problem has minimal business impact. The problem causes little impact on operations, or a reasonable circumvention to the problem was implemented.

Describing problems and gathering information

When describing a problem to IBM, be as specific as possible. Include all relevant background information so that IBM Software Support specialists can help you solve the problem efficiently. To save time, know the answers to these questions:

- Which software versions were you running when the problem occurred?
- Do you have logs, traces, and messages that are related to the problem symptoms? IBM Software Support is likely to ask for this information.
- Can you re-create the problem? If so, what steps were performed to re-create the problem?
- Did you make any changes to the system? For example, did you make changes to the hardware, operating system, networking software, and so on.
- Are you currently using a workaround for the problem? If so, be prepared to explain the workaround when you report the problem.

Submitting problems

You can submit problems to IBM Software Support in one of two ways:

Online Click Submit and track problems on the IBM Software Support site at http://www.ibm.com/software/support/probsub.html. Type your information into the appropriate problem submission form.

By phone For the phone number to call in your country, go to the contacts page of the *IBM Software Support Handbook* at http://techsupport.services.ibm.com/guides/contacts.html and click the name of your geographic region.

If the problem you submit is for a software defect or for missing or inaccurate documentation, IBM Software Support creates an Authorized Program Analysis Report (APAR). The APAR describes the problem in detail. Whenever possible, IBM Software Support provides a workaround that you can implement until the APAR is resolved and a fix is delivered. IBM publishes resolved APARs on the Software Support Web site daily, so that other users who experience the same problem can benefit from the same resolution.

Appendix J

Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents.

You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785 U.S.A.
For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement might not apply to you.

This informaiton could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner server as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758 U.S.A.
Such information may be available, subject to appropriate terms and conditions, including in some cases payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-levl systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for this specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similiarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using,

marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notices as follows:

©your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs.© Copyright IBM Corp. _enter the year or years_. All rights reserved.

If you are viewing this information in softcopy form, the photographs and color illustrations might not display.

Trademarks

IBM, the IBM logo, ibm.com, AIX, Candle, DB2, DB2 Connect, eServer, IBMLink, IMS, iSeries, IntelliWatch, Lotus, NetView, OS/390, Passport Advantage, ProductPac, pSeries, RACF, Rational, Redbooks, RMF, S/390, SAA, System z, SystemPac, Tivoli, the Tivoli logo, Tivoli Enterprise, Tivoli Enterprise Console, WebSphere, z9, z/OS, and zSeries are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Adobe, Acrobat, Portable Document Format (PDF), and Postscript are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.

Intel, Intel logo, Intel Inside®, Intel Inside logo, Intel Centrino®, Intel Centrino logo, Celeron®, Intel Xeon®, Intel SpeedStep®, Itanium®, and Pentium® are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a registered trademark of Linux Torvalds in the United States, other countries, or both.

Microsoft, Windows, and Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Oracle, JD Edwards, PeopleSoft, and Siebel are registered trademarks of Oracle Corporation and/or its affiliates.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, and service names may be trademarks or service marks of others.

Index

A

D

E

F

H

I

L

M

T

U

V

W