

IntelligenceCenter® Release Notes

Version 3.3.2.1

September, 2014

BLUE COAT

P/N 20-0335-3321 Rev. B

Disclaimer

THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY EXPRESS OR IMPLIED WARRANTY OF ANY KIND, INCLUDING WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT OF INTELLECTUAL PROPERTY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT SHALL BLUE COAT OR ITS SUPPLIERS BE LIABLE FOR ANY DAMAGES WHATSOEVER (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION, OR LOSS OF INFORMATION) ARISING OUT OF THE USE OF OR INABILITY TO USE THIS DOCUMENT, OR THE PRODUCTS DESCRIBED HEREIN, EVEN IF BLUE COAT HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. BECAUSE SOME JURISDICTIONS PROHIBIT THE EXCLUSION OR LIMITATION OF LIABILITY FOR CONSEQUENTIAL OR INCIDENTAL DAMAGES, THE ABOVE LIMITATION MAY NOT APPLY TO YOU. Blue Coat and its suppliers further do not warrant the accuracy or completeness of the information, text, graphics, links or other items contained within this document, or assume liability for any incidental, indirect, special or consequential damages in connection with the furnishing, performance, or use of this document. Blue Coat may make changes to this document, or to the products described herein, at any time without notice. Blue Coat makes no commitment to update this document.

Copyright/Trademarks/Patents

© 2014 Blue Coat Systems, Inc. All rights reserved. BLUE COAT, PROXYSG, PACKETSHAPER, CACHEFLOW, INTELLIGENCECENTER, CACHEOS, CACHEPULSE, CROSSBEAM, K9, DRTR, MACH5, PACKETWISE, POLICYCENTER, PROXYAV, PROXYCLIENT, SGOS, WEBPULSE, SOLERA NETWORKS, DEEPSEE, DS APPLIANCE, SEE EVERYTHING. KNOW EVERYTHING., SECURITY EMPOWERS BUSINESS, BLUETOUGH, the Blue Coat shield, K9, and Solera Networks logos and other Blue Coat logos are registered trademarks or trademarks of Blue Coat Systems, Inc. or its affiliates in the U.S. and certain other countries. This list may not be complete, and the absence of a trademark from this list does not mean it is not a trademark of Blue Coat or that Blue Coat has stopped using the trademark. All other trademarks mentioned in this document owned by third parties are the property of their respective owners. This document is for informational purposes only.

BLUE COAT MAKES NO WARRANTIES, EXPRESS, IMPLIED, OR STATUTORY, AS TO THE INFORMATION IN THIS DOCUMENT. BLUE COAT PRODUCTS, TECHNICAL SERVICES, AND ANY OTHER TECHNICAL DATA REFERENCED IN THIS DOCUMENT ARE SUBJECT TO U.S. EXPORT CONTROL AND SANCTIONS LAWS, REGULATIONS AND REQUIREMENTS, AND MAY BE SUBJECT TO EXPORT OR IMPORT REGULATIONS IN OTHER COUNTRIES. YOU AGREE TO COMPLY STRICTLY WITH THESE LAWS, REGULATIONS AND REQUIREMENTS, AND ACKNOWLEDGE THAT YOU HAVE THE RESPONSIBILITY TO OBTAIN ANY LICENSES, PERMITS OR OTHER APPROVALS THAT MAY BE REQUIRED IN ORDER TO EXPORT, RE-EXPORT, TRANSFER IN COUNTRY OR IMPORT AFTER DELIVERY TO YOU.

Americas:

Blue Coat Systems, Inc.

420 N. Mary Ave.
Sunnyvale, CA 94085

U.S. Government Restricted Rights

Rest of the World:

Blue Coat Systems International SARL

3a Route des Arsenaux
1700 Fribourg, Switzerland

Blue Coat software comprises "commercial computer software" and "commercial computer software documentation" as such terms are used in 48 C.F.R. 12.212 (SEPT 1995) and is provided to the United States Government (i) for acquisition by or on behalf of civilian agencies, consistent with the policy set forth in 48 C.F.R. 12.212; or (ii) for acquisition by or on behalf of units of the Department of Defense, consistent with the policies set forth in 48 C.F.R. 227-7202-1 (JUN 1995) and 227.7202-3 (JUN 1995). Blue Coat software is provided with "RESTRICTED RIGHTS." Use, duplication, or disclosure by the U.S. Government is subject to restrictions as set forth in FAR 52.227-14 and DFAR 252.227-7013 et seq. or their successors. Use of Blue Coat products or software by the U.S. Government constitutes acknowledgment of Blue Coat's proprietary rights in them and to the maximum extent possible under federal law, the U.S. Government shall be bound by the terms and conditions set forth in Blue Coat's end user agreement.

Revision History

July, 2014	IC 3.3.1.1
September, 2014	IC 3.3.2.1
October, 2014	Rev. B

Introduction

This document contains information specific to the IntelligenceCenter (IC) 3.3.x software only. If you are upgrading from an earlier version of IC, you can learn about other new features and software changes by consulting the release notes for the versions between your current software and v3.3.

See the following sections for specific information:

What's New in IC 3.3	page 4
Resolved Issues in IC 3.3	page 4
Fixed Security Advisory Issues	page 4
Interoperability Requirements	page 5
Preparing for Installation	page 6
Upgrading to IntelligenceCenter 3.3	page 8
Known Issues for IntelligenceCenter 3.3.1.1	page 16
Supplementary Information	page 26

What's New in IC 3.3

Support for PacketShaper S-Series

IC 3.3 includes support for the PacketShaper S500 running PS 11.1.1.14 or higher and PacketShaper S200 and S400 running PS 11.2.1 or higher.

Support for Windows 2012 Server

If you are installing IntelligenceCenter or DataCollector on a Windows 2012 Server, you need to run the installer in Windows 7 compatibility mode.

User Role Can View Scheduled Reports

Previously, scheduled reports were erroneously created as the system user, so only the Administrator role could view these reports. In IC 3.3, scheduled reports are now created as the user who scheduled the report. As a result, report visibility is now determined according to each users' group and role memberships.

Resolved Issues in IC 3.3

The following issues were resolved in **IC 3.3.2.1**:

- Fixed the issue in which the DataCollector failed to collect FDR data for certain PacketShapers due to incorrect decoding of the PacketShaper serial number. SR 2-970678202, 2-980546412, 2-985799006

The following issues were resolved in **IC 3.3.1.1**:

- Email addresses that are configured in scheduled reports are now displayed on the report configuration and scheduled tasks pages in the UI. SR 2-535909502
- Adding a syslog service no longer requires a device license.
- IC 3.3 now reports the correct PacketShaper link size for all versions of PacketWise. SR 2-517258652, 2-523376102SR, 2-517258652, 2-523376102
- IC is now able to successfully delete scheduled tasks. SR 2-508732142, 2-535909502, 2-754250884
- Backing up data to a CSV file failed because IC could not obtain the postgres password. This issue has been fixed. SR 2-505066562
- The reporting function `get_top_n_classes_by_group()` has been reworked to more properly filter leaf classes to avoid double counting class data. SR 2-458576522
- The class tree list is no longer lost after an IC upgrade. SR 2-423035231, 2-431663922, 2-437601252, 2-446090372.

Fixed Security Advisory Issues

The following Common Vulnerabilities and Exposures (CVE) have been addressed in IC 3.3:

- CVE-2012-2110,CVE-2012-2131
- CVE-2014-0224, CVE-2014-0221, CVE-2014-0195 & CVE-2014-3470.

Interoperability Requirements

Browser Requirements

IC has been tested with the following browser versions. You can log in to IC from any client machine that is equipped with one of the supported browsers.

- Google Chrome 35
- Firefox 30
- Internet Explorer 7.x, 8.x, or 9.x

Adobe Flash Player Requirements

Because the IC user interface is displayed using Adobe Flash Player, you must have a current version of Adobe Flash Player installed on the client system from which you will access IC. For the IC user interface to function properly, you must set the browser's Internet security settings to the default values. If you haven't already installed the latest version, make sure to do so before using IC. If you aren't sure which version of Adobe Flash Player is installed on your client system, go to:

<http://www.adobe.com/software/flash/about/>

To download the latest version, go to:

<http://www.adobe.com/products/flashplayer/>

If you do not have Flash installed and you attempt to log in to IntelligenceCenter, you will be redirected to the Flash download page.

PacketWise Requirements

IC can report data for the following Blue Coat appliances:

- PacketShaper or PacketShaper ISP appliances running PacketWise version 7.3.1 or higher
- PacketShaper S500 running PS 11.1.1.14 or higher

NTP Requirements

To ensure the integrity of the data collected by DC, you must configure it to use a Network Time Protocol (NTP) server. You will be required to provide the URL for an NTP server during the installation. For best results, DC and the PacketShaper appliances that it will use as data sources should all be configured to use the same NTP server. As a best practice, consider using the DC as the NTP server for your data sources. If the time between DC and the data sources drifts more than 2.5 minutes, DC may reject FDR records from the data sources, resulting in missing data or reports.

Preparing for Installation

Before you install this release, make sure that your system is prepared as follows:

- If you do not have a previous version of IC or DC installed, perform a fresh installation of the product using the procedures in the *IntelligenceCenter Getting Started Guide*. You can download this guide from <https://bto.bluecoat.com/documentation/pubs/view/IntelligenceCenter%203.3>
- Upgrade to 3.3 is supported from IntelligenceCenter 3.2. For instructions, see “Upgrading to IntelligenceCenter 3.3” on page 8.
- Whether you are upgrading or performing a new installation, make sure your server(s) meet the system requirements appropriate for the deployment you are planning. The system requirements are documented in the *IntelligenceCenter Getting Started Guide*.
- You must log in to the server where you plan to install IC and/or DC using a user account with administrative rights that include the following:
 - Read/Write permissions on the file system
 - Permissions to access and modify registry
 - Permissions to restart the server
 - Permissions to deploy/install software
 - Permissions to "Log on as a service" and "Log on locally" (SR# 2-285642996)
- In environments where Windows Group Policy Objects (GPOs) control access to user and computer accounts, you may have to use the local admin account rather than a domain account within the Administrator's group to ensure that you have sufficient rights to install IC.
- You must install IC and DC on a server that is running an English Language version of Microsoft Windows Server. In addition, the Microsoft Windows Regional and Language option on the server(s) where you are installing IC and/or DC must be set to English and remain set to English after installation to ensure proper reporting. After installation, you can, however, change the Regional and Language options Text services and input languages on the Language tab to use a non-English keyboard. For additional information about system requirements, refer to the *IntelligenceCenter Getting Started Guide*.
- You must install IC and DC on a dedicated system. In addition, you must make sure the following services are disabled before installing the software:
 - Anti-virus programs
 - Windows update services
 - Backup agentsNote that you can re-enable anti-virus programs and backup agents after installation. However, to prevent degradation in performance, you must exclude the Postgres data directories from scanning/backup.
- PostgreSQL requires the following file permissions for a successful installation:
 - Postgres needs read/execute rights to the C:\ drive. Use either the Users or Authenticated Users (preferred; more secure) with read/execute rights.
 - Postgres needs read/write/execute rights to the data location (C:\ProgramFiles\PostgreSQL\8.4\data by default). Use either the Users or Authenticated Users (preferred; more secure) with read/execute rights.
- The local PostgreSQL service account must have rights to Log on as a service and Log on locally.
- If you previously tried to install or uninstall IntelligenceCenter or DataCollector and the process failed to complete successfully, you must clean up the Windows registry before attempting a re-install. Run one or both of the following batch files to clean up the registry, depending on which process failed:

-
- If the IC install or uninstall failed, run the `cleanup_registry.bat` file, which is located in the `\bin` folder in the DataCollector installation location (C:\BlueCoat\IntelligenceCenter by default).
 - If the DC install or uninstall failed, run the `cleanup_registry.bat` file, which is located in the `\bin` folder in the DataCollector installation location (C:\BlueCoat\DataCollector by default).

Note that although these scripts have the same name, they are different. You should only run the script in the IC installation folder to clean up IC and you should only run the script in the DC installation folder to clean up DC.

- For more detailed information on how to prepare your system for an IC installation, refer to KB article KB3932 (<https://kb.bluecoat.com/index?page=content&id=KB3932>).
- If you previously tried to migrate from the previous version of IntelligenceCenter or DataCollector to the latest version of IntelligenceCenter or DataCollector and the process failed to complete successfully, you need to re-install IntelligenceCenter or DataCollector and restore the backup. Before you re-install and try to restore a backup, you must clean up the postgres. During the postgres cleanup, the “postgres” windows user also must be deleted.
 - To delete a user, perform the following steps:
 - Select **Start > Control Panel** and click **User Accounts**.
 - Select the postgres user.
 - Select **Delete the account**.
 - At the confirmation box, select **Delete Files** to remove postgres user files.




Note: Ensure that all the files associated with the user are removed completely from the Windows server. Failure to do so may result in PostgreSQL installation failure and IntelligenceCenter/DataCollector installation failure as well.

Upgrading to IntelligenceCenter 3.3

You can upgrade from IC 3.2 to 3.3. When you upgrade, all of your data, configuration settings, and log files will be preserved. Keep in mind that in order for IC and DC to communicate, you must upgrade both IC and DC so that they are both running the same version of the software. For example, you cannot run a 3.3 DC with a 3.2 IC Server.

When you perform an upgrade, the database schema is modified and all of your existing data is migrated into the new database. Depending on the amount of data you have and the processing power on the system where you are running IC and/or DC, this migration can take many hours and you will not be able to use IC during this time (when you log in, a message will indicate that migration is in progress). In addition, it can take several hours for the current database tables to catch up with any new data that is being collected post-migration.

 **Note:** If you do not need to upgrade (that is, you do not want to keep the data you have previously collected and the configuration settings you have defined), you should perform a fresh installation. To do this, you must uninstall the existing version of IC and DC and then install the 3.3 software. For instructions, refer to the *IntelligenceCenter Getting Started Guide*.

In addition, although the IC and DC upgrade process automatically migrates your existing databases, Blue Coat strongly recommends backing up your existing IC and DC databases before performing an upgrade.

The following sections describe how to back up your existing IC and DC installations and then upgrade each component:


- “Backing Up the Databases” on page 9
- “Upgrading DataCollector” on page 9
- “Upgrading IntelligenceCenter” on page 12
- “Connecting to PostgreSQL 8.4” on page 15
- “Verifying the Upgrade” on page 15

Backing Up the Databases

Although the upgrade to IC 3.3 migrates your existing IC and DC databases, you should back up the databases before upgrading for added security.

The backup and restore scripts are installed to the `C:\Program Files\PostgreSQL\8.4\tools` folder on the server where IC is installed.

The following procedure describes how to back up a 3.x database.

 **Note:** If IC and DC are running on the same server, this script will automatically back up both databases. If IC and DC are running on separate servers, you must run the script on both systems. You **MUST** back up both components in order to be able to successfully restore them later.

1. Open a Windows command shell.
2. Create the directory where you want to store the backup files. As an extra security measure, you may want to create the backup directory on a different system and/or disk drive than your IC and DC databases are installed on in case there is a system failure.
3. Change to the `C:\Program Files\PostgreSQL\8.4\tools` directory.
4. At the command prompt enter the following command:

```
python pg_backup.py -b <backup_location>
```


where `<backup_location>` is the path and name of the folder you created to store the backup files. For example:


```
python pg_backup.py -b H:\Database_backups
```

When the script finishes, the command prompt returns.
5. Go to the directory you specified when you ran the script. You will find the following backup files:
 - `db_collector_yyyymmdd.dmp`—This is the backup for the DataCollector database. The `yyyymmdd` represents the date that the backup was created. For example, if you ran the script on February 3, 2014, the file name would be `db_collector_20140203.dmp`. This file will only be present if DC is installed on the system.
 - `db_PolicyVisionDB_yyyymmdd.dmp`—This is the backup for the IntelligenceCenter database. The `yyyymmdd` represents the date that the backup was created. For example, if you ran the script on February 3, 2014, the file name would be `db_PolicyVisionDB_20140203.dmp`. This file will only be present if IC is installed on the system.
6. If IC and DC are not installed on the same server, repeat this procedure on the server where the other component is installed. To ensure that the databases stay in sync, the backup files for IC and DC must be created at roughly the same time. If the components are installed on separate systems, make sure you run your backups at the same time or one after the other.

Upgrading DataCollector

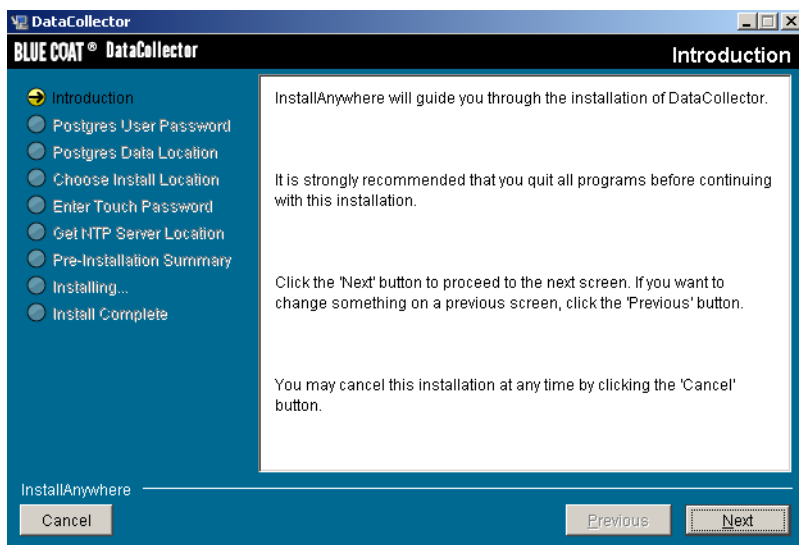
If you have a DC version 3.2 installed, you can upgrade to the current version without uninstalling the previous version. When you upgrade, all of your data, configuration settings, and log files will be preserved. In addition, the setup will automatically use the same preferences, such as installation and data location and NTP server, that you selected during the original installation.

 **Note:** Keep in mind that in order for IC and DC to communicate, you must upgrade both components. For instructions on upgrading IC, see “Upgrading IntelligenceCenter” on page 12. If you are running IC and DC on the same system, **DO NOT** reboot the system until you have upgraded both IC and DC.

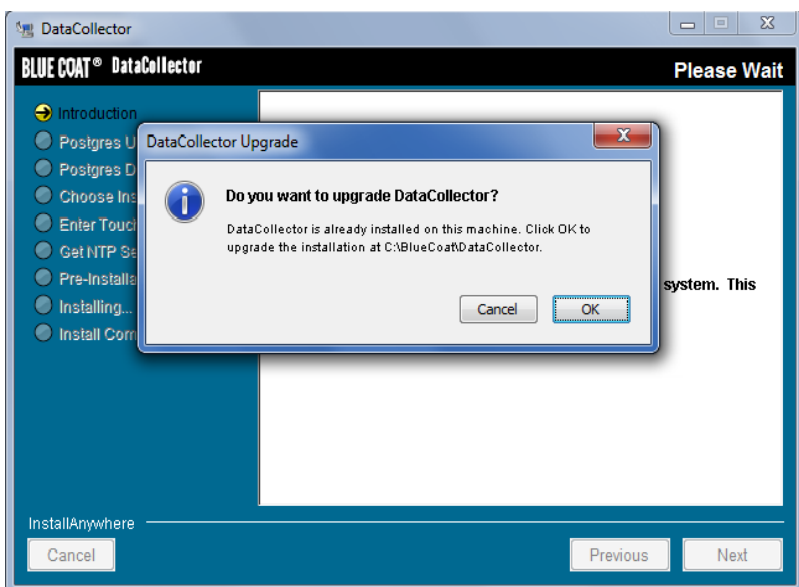
 **Note:** Although the upgrade migrates your existing database schemas, you may want to back up your database before proceeding with the upgrade. See “Backing Up the Databases” on page 9 for more information.

To upgrade DC, quit all open programs—including anti-virus software—and then do the following:

1. Log in to the system where you plan to install DC using an account with administrative privileges. You will not be able to successfully install DC if you do not have administrative rights.
2. Download the *DataCollector_installer.exe* file and save it to the hard drive on the system where you plan to install DC.
3. Double-click the *DataCollector_installer.exe* file. The *Introduction* screen appears.

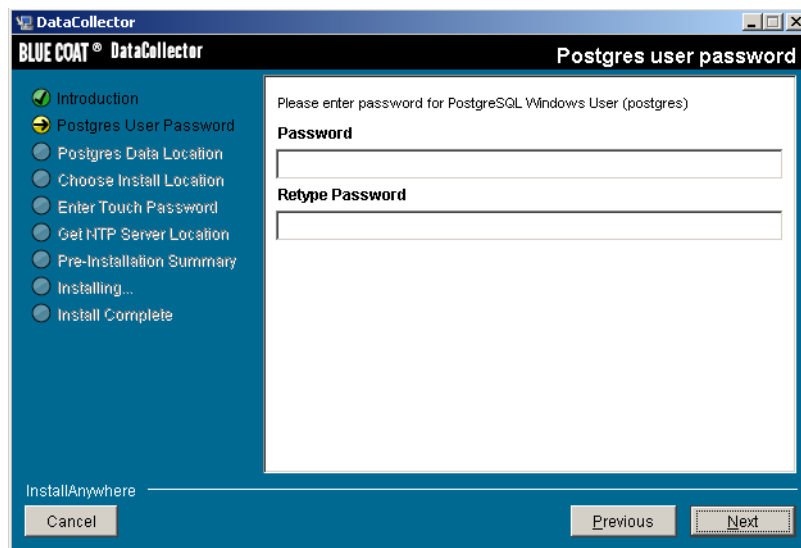


4. At the *Introduction* screen, click **Next**. The *Upgrade* dialog box is displayed. Keep in mind that when you upgrade, the data migration process can take several hours to complete after you complete the software installation. If you don't need to keep your data, consider uninstalling the previous version of the software and then doing a new install. However, if you plan to upgrade IC, you must also upgrade DC.

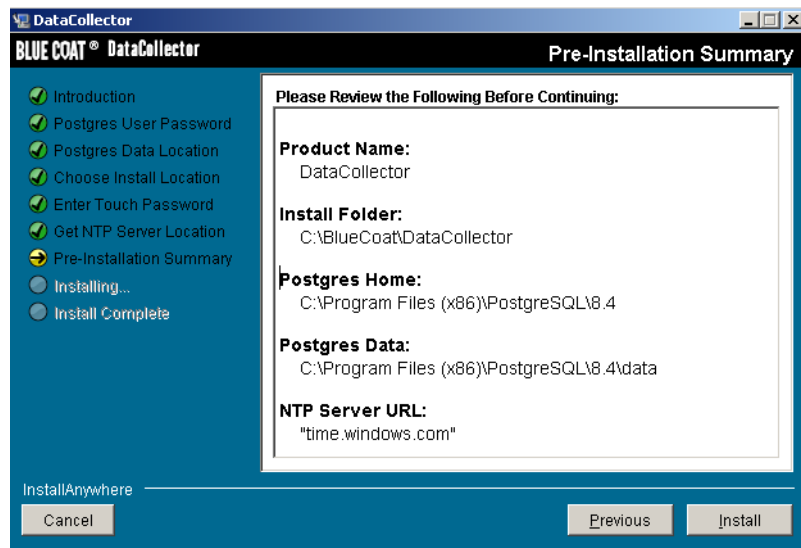


5. If you want to upgrade, click OK. Click **Next** to confirm that you want to proceed with the upgrade. The previous version of the software is uninstalled.

6. At the *Postgres user password* screen (shown below), create a password. For verification purposes, reenter the password in the **Retype Password** field. Click **Next**.



7. Review the *Pre-Installation Summary* screen, and then click **Install**. Notice that the setup uses the same installation settings you selected when you originally installed the product..





During installation the *Installing DataCollector* screen shows the progress of the installation. When the installation completes, the *Install Complete* screen is displayed.

8. Before you complete the setup, use a text editor to open the `\jboss-4.2.2\conf\wrapper.conf` file, which is located in the DC installation folder (`C:\BlueCoat\DataCollector` by default) and change the value of the `host` parameter from `0.0.0.0` to the actual IP address of the system where you just installed DataCollector.
9. Go back to the *Install Complete* screen and select whether to reboot the system now or later and then click **Done** to complete the upgrade. If you are installing IC and DC on the same system, you must install both components before rebooting. If you have not yet installed IC, continue to the next section.

Upgrading IntelligenceCenter

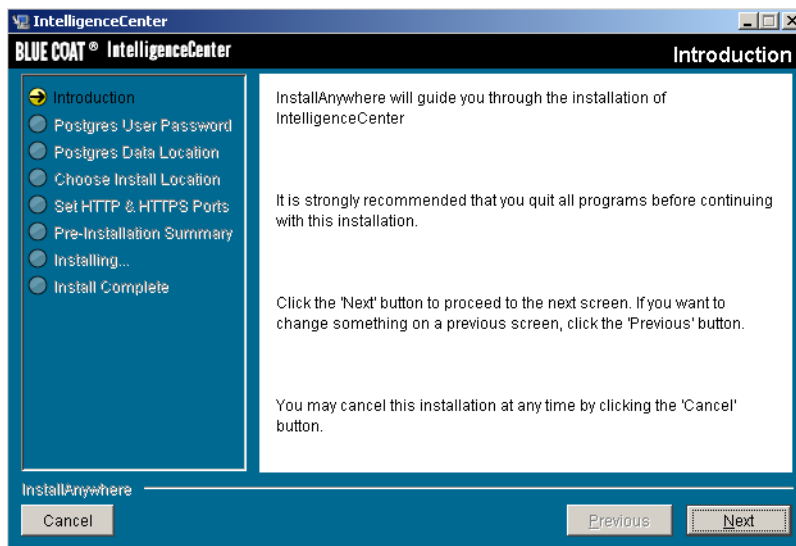
If you have a 3.2 version installed, you can upgrade to the current version without uninstalling the previous version. When you upgrade, all of your reports, configuration settings, and log files will be preserved. In addition, the setup will automatically use the same preferences, such as installation and data location and HTTP port number, that you selected during the original installation.

 **Note:** Keep in mind that in order for IC and DC to communicate, you must upgrade both components. For instructions on upgrading DC, see “Upgrading DataCollector” on page 9. If you are running IC and DC on the same system, DO NOT reboot the system until you have upgraded both components or they will no longer be able to communicate.


 **Note:** Although the upgrade migrates your existing database schemas, you may want to back up your database before proceeding with the upgrade. See “Backing Up the Databases” on page 9 for more information.

To upgrade IC, quit all open programs—including anti-virus software—and then do the following:

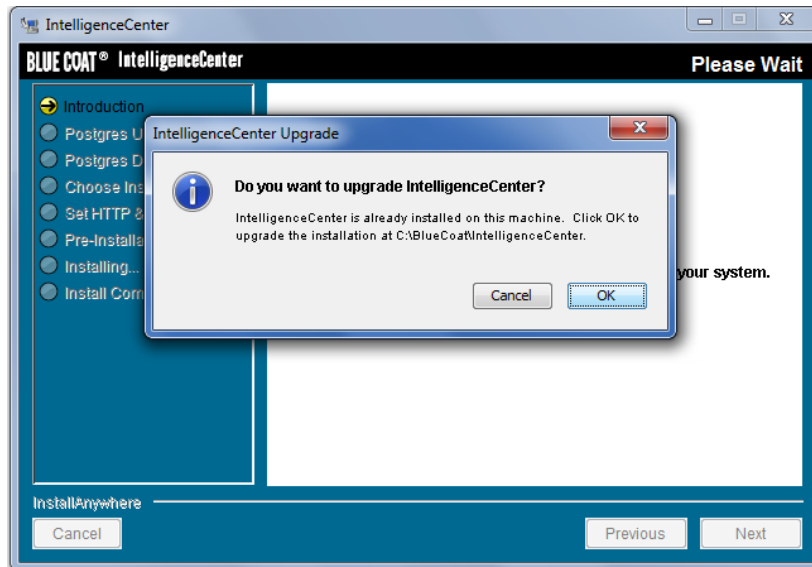
1. Log in to the system where you plan to install IC using an account with administrative privileges. You will not be able to successfully install IC if you do not have administrative rights.
2. Download the *IntelligenceCenter_Installer.exe* file and save it to the hard drive on the system where you plan to install IC.
3. Double-click the *IntelligenceCenter_Installer.exe* file. The *Introduction* screen appears.



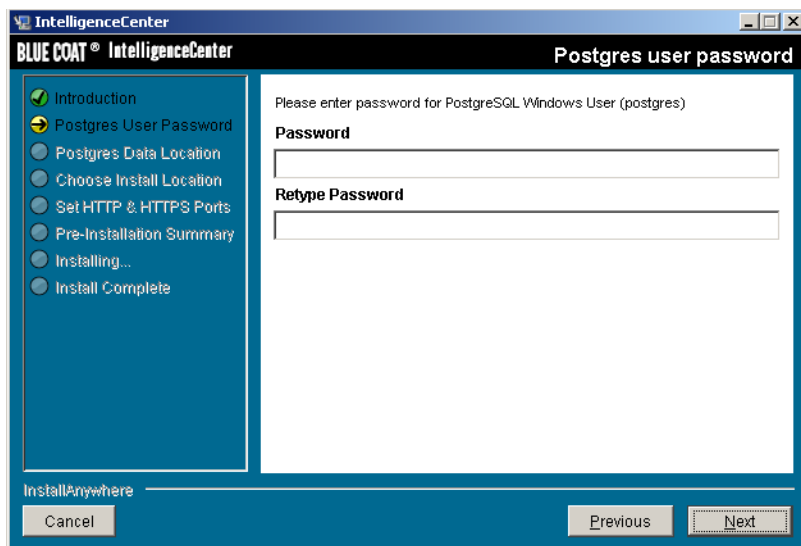
4. At the *Introduction* screen (shown above), click **Next**. The *IntelligenceCenter Upgrade* dialog box is displayed. If you want to upgrade, click OK. Click **Next** to confirm that you want to proceed with the upgrade. The previous version of the software is uninstalled

 **Note:** Keep in mind that when you upgrade, the data migration process can take several hours to complete after you complete the software installation. If you don't need to keep your data, consider uninstalling the

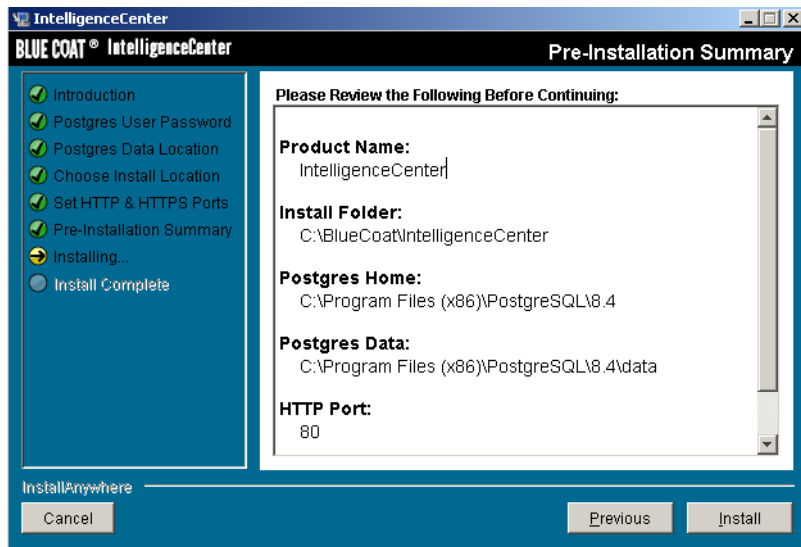
previous version of the software and then doing a new install. However, if you plan to upgrade DC, you must also upgrade IC.



5. At the *Postgres user password* screen, create a password. For verification purposes, reenter the password in the **Retype Password** field. Click **Next**.



6. Review the *Pre-Installation Summary* screen, and then click **Install**. Notice that the setup uses the same installation settings you selected when you originally installed the product. During installation the *Installing IntelligenceCenter* screen shows the progress of the installation.



7. Select whether to reboot the system now or later and then click **Done** to complete the upgrade. If you are installing IC and DC on the same system, you must install both components before rebooting. If you have not yet installed DC, go to "Upgrading DataCollector" on page 9.

Connecting to PostgreSQL 8.4

Follow these steps to connect to PostgreSQL 8.4 for the first time:

1. After installing IC and DC, you need to get the PostgreSQL password. Go to the following registry location:
`HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\PostgreSQL\BlueCoatIntelligenceCenter`
 2. Get the dbPassword value.
 3. To decode the hex value, use this website: <http://www.base64decode.org/>
 4. Open pgAdmin3.
 5. Right-click **PostgreSQL8.4 (localhost:xxxx) database**
 6. Click **connect**.
 7. When prompted for the password, enter the decoded password.
- You are now connected to PostgreSQL 8.4.

Verifying the Upgrade

After you upgrade IC and DC and reboot the server(s), you must verify that the upgrade completed successfully. To verify the upgrade:

1. Log in to IC by going to one of the following URLs:
`https://<ip_address>:<port>` (for a secure connection)
- or -
`http://<ip_address>:<port>` (for a non-secure connection)

where <ip_address> is the address of the server where you installed IC. You only need to include the <port> if you specified a port number other than the default when you installed IC.
2. When prompted, enter the following login credentials and then click **Login**:

User Name: admin
Password: admin
3. Select **Configure > Network**.
4. Locate the DataCollector object in the list of network devices in the left pane and make sure that the corresponding status icon is green.
5. Select the DataCollector object. The **Status** tab opens in the right pane.
6. To verify that the upgrade was successful, make sure the **Collection Status** field shows **Collecting** (green).
7. If either the DC object is red or the Collection Status shows Not Collecting (red), you must close the IC application and restart the following Windows services:
 - Blue Coat DataCollector 3.3.1.1
 - Blue Coat Agent For DataCollector 3.3.1.1

Known Issues for IntelligenceCenter 3.3.1.1

The following sections list issues discovered in the IC and DC software:

- “Install/Upgrade/Uninstall Issues” on page 16
- “Device Management Issues” on page 17
- “IC Interface Issues” on page 18
- “Restart the Blue Coat IntelligenceCenter 3.3.1.1 service.” on page 18
- “Compatibility Issues” on page 20
- “Authentication Issues” on page 20
- “Data Collection Issues” on page 20
- “User Issues” on page 21
- “Portlet Issues” on page 21
- “Reporting Issues” on page 22

Install/Upgrade/Uninstall Issues

- If you have problems launching the software after installation, check the following PostgreSQL settings:
 - Make sure the Postgres Windows user is a member of the Users group, but NOT a member of the Power Users or Administrator groups.
 - Make sure the Postgres user has permissions to Log on Locally, as well as Read, Update, and Execute permissions on the %PG_HOME% and %PG_DATA% directories (and apply these permissions recursively to all sub-folders). To do this:
 1. Right-click the PostgreSQL folder created by the installation and select **Properties**.
 2. On the **Security** tab, click **Edit > Add** and enter **postgres** in the **Enter the object names to select box**.
 3. Click **Ok**.
 4. In the Group or Usernames dialog box, select the Postgres account. Select the **Full Control** checkbox in the **Allow** column next to the Postgres account (or, select **read/write** if you do not want to give full permissions).
 5. Assign **Logon as a service** permissions to the account.
- The PostgreSQL installation contains `libeay32.dll` and `ssleay32.dll` files. However, if you already have these files on the system where you are installing, the PostgreSQL installation will not overwrite them. If you have older versions of these files, rename them so that the PostgreSQL installation will install the version it needs.
- The following error message is included in the IntelligenceCenter `server.log` file after installation:

```
org.eclipse.birt.report.model.metadata.DimensionPropertyType  
validateUnits SEVERE: unit:% not allowed
```

This message is normal and does not indicate a problem with the installation.

- Although you can install IC and/or DC using a remote desktop connection, you must launch remote desktop with the following options:

```
mstsc.exe -v <ip-address> -console
```

This is a bug in the PostgreSQL installer; for details, go to:

http://www.postgresql.org/docs/faqs.FAQ_windows.html#3.5

- Sometimes when installing DC on a 64-bit system, the installer cannot get the IP address from the system and sets the DC address to 0.0.0.0. Then, when you try to add the DC to the IC network using the IP address of the machine, IC will be unable to connect to DC as shown in the following screen shot:

In this case, you can add DC to the IC network by editing the `\jboss-4.2.2\conf\wrapper.conf` file, which is located in the DC installation folder (C:\BlueCoat\DataCollector by default) to reflect the actual IP address of the system.

Device Management Issues

- **Devices**—When you add a new device, IC does not require that the IP address is valid and reachable or that the serial number and/or password are valid. However, these fields must be correct in order for you to report on the device or use the single sign-on feature.
- **Devices**—When adding a PacketShaper to the IC network, make sure you use the touch password, not the look password. If you use the look password, IC will not be able to get the correct time zone information from the appliance.
- **Devices**—IC retrieves statistical information (such as link size, traffic discovery, acceleration, and compression) about the devices it manages when you add the device to the IC network topology. However, it does not automatically refresh this information after the initial import. To force IC to retrieve updated statistics, click the **Sync** button on the **Configure > Devices** page (keep in mind, however, that if you have a large number of devices, the update operation will degrade system performance for up to 30 minutes). Note that the **Last Contact** field indicates the date and time that IC last successfully pinged the device, not the time of the last update.

-
- **PolicyCenter Import**—Although the PolicyCenter import feature of IC allows you to quickly and easily import devices into IC, it should only be used when your PolicyCenter groups use unique names. If PolicyCenter groups do not have unique names, IC will consolidate the groups upon import, resulting in unwanted movement in complicated layouts. The only way to prevent this is to either manually import your devices into IC or to rename your groups in PolicyCenter so that they are unique before using the IntelligenceCenter import feature. For more details on this issue, refer to Knowledge Base article TFA34 (<https://kb.bluecoat.com/index?page=content&id=TFA34>).

IC Interface Issues

- **Applications**—When defining matching rules for an application, IC matches complete traffic class names only; it does not perform string matching. For example, if you want to define matching rules for your discovered TCP ports (TCP_Ports_124, TCP_Ports_125, TCP_Ports_4235, etc.), you must define a separate matching rule for each traffic class; you cannot use a wildcard string to match the TCP_Ports portion of the traffic class name.
- **Applications**—When defining **Included Traffic Classes** for an application, the /Inbound and/or /Outbound are automatically prepended to the path as directional matching attributes (match inbound only, outbound only, or both):
 - **Inbound Only**—Only classes under the /Inbound branch will be matched against.
 - **Outbound Only**—Only classes under the /Outbound branch will be matched against.
 - **Both Directions**—All classes are matched against.

Therefore, you do not need to explicitly define /Inbound and /Outbound in your application definitions.

- **Class Discovery**—Although the IC interface provides an option to filter classes based on whether they are active or inactive, these filters are currently not functional. If you select Inactive, the class list will be empty.
- **Role Based Access Control**—By default, the reports you can run are arranged hierarchically by report type (Application, Device, Host, Site). However, if you create role-based access rules that limit which reports a particular role can view or run, the report hierarchy will no longer be displayed.
- **Role Based Access Control**—If you configure a user to receive email notification of alerts and the user is assigned to multiple roles, the user will receive multiple email notifications for each alert (one for each assigned role).
- **Traffic Classes**—If a traffic class is deleted in PacketShaper, the class is still displayed when you select the **Active** radio button in the IC user interface. All updated classes are displayed in IC, but the old classes are still displayed as active classes in IC.
- **View Reports**—By default, the **Archived Reports** pane can display a maximum of 300 archived reports. If you archive a large number of reports, you may not be able to see all of the reports that have been generated; instead you will see the first 300 reports in the selected display range. By default, the **Archived Reports** pane shows the first 300 reports archived **This Week**. You can filter what reports are displayed by selecting a smaller time range in the **From** drop-down list or you can filter based on a specific time range and matching string (select **Other** in the **From** drop-down list). You can also manually edit the `ic-system-configuration.xml` file to increase the maximum number of archived reports that display (up to a maximum of 500). To modify the setting:
 - Open the `ic-system-configuration.xml` file in a text editor. This file is located in the `\apache-tomcat-6.0.18\webapps\ROOT\conf` directory in the IntelligenceCenter installation location (C:\BlueCoat\IntelligenceCenter by default).
 - Locate the `<max-reports-loading>` parameter and edit the value to the new maximum number of archived reports you want to be able to display (up to a maximum of 500).
 - Save the XML file.
 - Restart the Blue Coat IntelligenceCenter 3.3.1.1 service.

Browser Issues

- **Browser Controls**—IC does not support the browser Back and Forward buttons. Attempts to use these buttons from within IC will log you out of the system. Instead, use the IC menus to navigate from screen to screen.
- **Browser Controls**—Attempts to refresh your session using the browser refresh button will log you out of the system.
- **Browser Display**—Font smoothing is not enabled by default on Windows XP. To improve font display on Windows XP systems, you can enable ClearType:
 1. Right-click the Windows desktop and choose **Properties**.
 2. Click the **Appearance** tab.
 3. Click the **Effects** button.
 4. Select the checkbox **Use the following method to smooth edges of screen fonts** and then choose **ClearType** from the drop-down list.
 5. Click **OK**.
- **Firefox**—The Adobe Flash player plug-in for the Firefox browser does not allow file uploads over HTTPS when the web server uses a self-signed certificate. Therefore, you cannot use Firefox to upload files (such as licenses) to IC over HTTPS. You can upload files over HTTP or you can use HTTPS with a valid certificate. For more information about this Adobe issue, go to:
<http://bugs.adobe.com/jira/browse/FP-226>
- **HTTPS**—If you are using HTTPS with the default certificate provided with IC, the browser may refuse to communicate with the server and the application will hang. If you want to use HTTPS, you must install a new SSL certificate (purchased from a Certificate Authority) on the IC server. Refer to the *IntelligenceCenter Getting Started Guide* for instructions on installing a new certificate.
- **Internet Explorer**—When using IE8 to run IC reports remotely, you cannot open reports in the Microsoft Word (.doc) or Microsoft PowerPoint (.ppt) formats due to the following known Microsoft issue:
<http://support.microsoft.com/default.aspx?scid=http://support.microsoft.com:80/support/kb/articles/q316/4/31.asp&NoWebContent=1>
As a workaround, use Firefox to open these types of report formats.
- **Internet Explorer**—When using IE, if you configure and then minimize four or more portlets and then use one of the View Options to arrange them (Tile Portlets, Tile Portlets Vertically, Tile Portlets Horizontally, or Cascade), IE will return an exception.
- **Internet Explorer**—When a user who is assigned a role that does not include permission to see network views clicks on the Views tab when selecting a Network Group or Device to which to restrict a report an error occurs. This issue only occurs with some versions of IE 7, such as 7.0.6002.18005 (256-bit encryption).
- **Internet Explorer**—You cannot open two separate IC logins from the same IE 8 browser because IE 8 does not create a separate process for each session. As a workaround, manually start a new session in IE before attempting to log in to IC under a second account:
 1. In IE 8, launch IC and log in using the credentials for the first user account.
 2. Select **File > New Session**. IE launches a new browser window.
 3. Launch IC in the new browser window and log in using the credentials for the second user account.

Compatibility Issues

- Microsoft Java VM security update 816093 may conflict with IC. Do not install this update on the server where IC and/or DC are installed.
- IC and PolicyCenter both use 80 as the default port for HTTP, which causes a port number conflict if these applications are installed on the same server. By installing IC on a dedicated server as recommended, you can prevent this type of port number conflict and ensure that IC has the system resources it requires.

Authentication Issues

- The first time IC successfully authenticates a user using the external LDAP service, it creates an IC user account for the user. However, IC will not log the user in upon this first successful authentication. Instead, it displays a message on the login screen indicating that the administrator must activate the account. In addition, IC should generate an alert account indicating that the account requires activation. However, the alert that gets created incorrectly states that the login failed due to an incorrect password.

Data Collection Issues

- DC gets some of its PacketShaper-specific information from ME data. However, if you plan to collect FDR data but not ME data, DC may not be able to get all of the information it needs. In this case, you must configure your PacketShaper data sources to emit Packeteer-0 data to DC in addition to Packeteer-2 data. Packeteer-0 packets are mapping messages that allow DC to decipher PacketShaper-related information in the FDRs they receive. For example, in the FDR's ClassID field, an ID identifies the traffic class. In order for the collector to understand what class is actually associated with the ID, it uses the class map in the Packeteer-0 record. This class map lists each traffic class on the unit along with the identifying number assigned to each class.

To emit Packeteer-0 to all collectors defined for a PacketShaper appliance, go to the command line interface and enter the following command:

```
setup variable flowRecordsSendPktr0 1
```

- If the PacketShaper appliance from which you are collecting ME data become overloaded, DC may stop collecting from them (although class discovery will continue). To determine whether DC has stopped collecting ME data, go to the DataCollector Status tab (**Organize > Devices > DataCollector > Status**). In the *Metric Data Collection Status* section of the screen look to determine the date and time of the **Last data retrieval** (and whether the time is close to the time of the **Last type/class/partition discovery**). Normally the **Last data retrieval** time and the **Last type/class/partition discovery** time would be close together. If the values are not close together and if the **Last data retrieval** was more than two hours ago (as is the screen shot below), ME collection has stopped.

Metric Data Collection Status	
Last type/class/partition discovery	4:32 PM
Last data retrieval	Nov 23, 2008 8:01 AM

To resume ME collection, restart the Blue Coat DataCollector 3.3.1.1 service.

- In some cases, the class discovery process can get behind if you have a large number of slow PacketShaper appliances configured as data sources. If this happens, ME data collection will also fall behind (see the description in the issue above to determine if ME data collection is behind or stopped). To work around this problem, change your **Query interval** (on the DataCollector **Configuration** tab) to 60 minutes and the **Data granularity** to 15 minutes.
- DataCollector cannot collect data from a PacketShaper that is configured with one of the following time zones due to a defect in the PacketWise XML API:

- GMT -12:00 Baker Island, Howland Island
- GMT +00:00 Local time, pending time zone configuration
- GMT +00:00 Casablanca, Monrovia


The DataCollector **Status** tab will show the following error for Shapers with one of these time zones configured: The selected time span is not valid. Either the DataCollector and ME data source clocks are not synchronized or the data source time zone has changed.

- If you have set up FDR data collection on DC and later want to stop it or change the port number on which DC listens for FDR, you must first stop all of your data sources from emitting FDR to DC and then make your change. If you stop FDR data collection on DC (by unchecking the **Collect Flows** checkbox on all of your data sources) and/or change the FDR port number on DC (using the **Data port** field on the DC **Configuration** tab) before stopping the data sources from emitting FDR to DC, DC will shut down its FDRListener, which could cause it to hang. In this case, you must restart the Blue Coat Agent For DataCollector 3.3.1.1 and Blue Coat DataCollector 3.3.1.1 services.
- DC does not regulate FDR traffic and could, therefore, become overloaded during periods of high volume FDR traffic.

User Issues

- Occasionally, users are unable to see alerts even though their assigned role allows it.
- When a user modifies an object, other users with concurrent IC sessions will not be able to see the changes. For example, if one user deletes an object — such as a device entry or a network view — other users who are concurrently using IC will still see the object. To protect the IC database from becoming corrupted, attempts to select or edit the object will then result in a server error. Similarly, if another user creates a new object, such as a view or a site, other users with concurrent IC sessions will not be able to see the new object. In order to see the current state, users must refresh their IC sessions.
- Each person who logs in to IC should have a unique user name and password. IC maintains settings (such as single sign-on credentials and portlet configurations) for each user profile. If multiple users log in using the same user profile concurrently, changes to the profile-specific configurations can conflict.

Portlet Issues

- **Class Utilization Portlet**—Sometimes when viewing stack charts the browser will crash when you move the sliders to change the time range displayed on the graph due to an Adobe Flash issue. This issue does not occur when viewing line charts. To resolve the issue, restart the browser and log back in to IC.
- **Timezone Offset**—If you are running IC from a client system that is set to use a time zone with an offset \pm HH:30 or \pm HH:45, the portlets will display data starting 12:30AM or 12:45AM rather than 12:00AM when you select the Today or Yesterday time range. There is no workaround for this issue.
- **Top N Children portlet**—The unit for 'Partition Burst Limit' should be interpreted as bits per second (bps) even though it is displayed as "B."
- **VoIP Performance portlet**—When you select a section on the time-selector bar, the shaded area on the graph sometimes does not align with the corresponding x-axis labels. Note that the data displayed in the tables below the time-selector bar correspond to the time period selected on the time-selector bar only.
- **VoIP Performance portlet**—If you make a change to the configuration and then cancel the configuration without saving your changes, the portlet no longer displays properly. To resolve this issue, click the  icon and then select **Refresh Data**.

- **Top N Children portlet**—If you configure the Max Count to 20 or greater, the tick marks on the pie chart do not always display properly. To work around this issue, either reduce the Max Count or increase the size of the portlet.
- **Multiple Portlet Instances**—If you are running IC in an environment with a high number of flows (in excess of 500,000 flows per minute) and you are running many portlet instances, the IC GUI may become unresponsive and DC may crash.

Reporting Issues

- **Applications**—The pre-defined Oracle application groups together Oracle database, client, and management traffic. However, this may not be useful from an application reporting perspective. To change this, you can manually modify the Oracle application definition or define your own custom applications for reporting on the Oracle applications you are interested in.
- **Applications**—Some reports, such as the Top Talkers and Top Listeners reports, provide links to the Top Applications report. However, application traffic will not be displayed in an IC report unless there is an application defined in IC to represent it. If the traffic reported in the Top Talkers and Top Listeners reports do not have corresponding applications defined in IC, the corresponding traffic will be aggregated into an application called “Undefined.” However, because this “Undefined” application really represents a grouping of traffic that may or may not be related, you cannot drill-down into the data or into time-series charts as you can with other applications.
- **Business Hours Reporting**—In DC, all data gets stored in tables in GMT to allow for normalization of data for reporting in different time zones. Therefore, when you configure business hours reporting, the start and end times that you set (which are based on the DC time zone) get converted to GMT within the database. For example, suppose your DC is set to use time zone -0800 and you configure business hours to start at 2100 and end at 0600, the configuration gets converted to GMT with a start time of 0500 and an end time of 1400. Thus, because the configured business hours occur late in the day, some hourly data that actually belongs to one day will be stored in tables for the next day (based on GMT). Due to differences in how and when data is rolled from one table to the next, this can cause some reporting anomalies. For example, all hourly data that is beyond the GMT date line will be stored in the hourly table for the next day. At the end of the next day, DC will make the business hour time zone adjustments and store the data in the proper daily table. However, rollups of day data to the monthly table and monthly table to the yearly table occur at the end of the GMT day on the last day of the month and therefore the extra hourly data that is initially stored in the table for the next day will not get rolled up into these tables. This can cause some reporting discrepancies across different time ranges, depending on the date and the tables from which DC must fetch data to render the report. Here are some examples:
 - Reports with a time range of Last Month or Last Quarter will be missing a day because the monthly table rollup and the yearly table rollup occurs at the end of the last day of the month based on GMT. However, data for the last day of the month will be rolled up into the next day (the first day of the next month). Therefore, when the rollup for the month occurs (based on GMT), the data the last day of the month will not be included.
 - Reports for the same data but different time ranges may show different information, depending on which database tables the data comes from and the date on which you run the report. For example, reports for Last Month and Last Quarter will have a one-day discrepancy when run after the first of the month because the Monthly data rolled up into the Yearly table without taking into account data for the last day of the month that occurred in the next GMT day. However, after the first day of the month, the hourly data that was initially stored in the table for first will be rolled to the proper Daily table for the last day of the month and then to the proper Monthly table.
- **Business Hours Reporting**—If you configure business hours for reporting and later change them, IC reports will display an overlap in the defined business hours for a short period of time while it adjusts to the new settings to ensure availability of data.

-
- **Business Hours Reporting**—If you run a report that includes a time series plot using a time range greater than one month, the timestamp for the plot will always be displayed as starting at 00:00 (midnight). However, if you have business hours reporting enabled, the actual start range for the plot will be the start of your business hours even though it will be displayed as 00:00. For example, if your business hours are set to 8am to 6pm, the time series will show that it is reporting data starting at 00:00 when in actuality the data that is displayed starts at 08:00.
 - **Business Hours Reporting**—When you first configure business hours reporting, the time zone automatically adjusts if daylight saving time (DST) is in effect. However, if there is a switch between DST and standard time after the initial configuration of the feature, the configuration does not automatically change. As a workaround, you must go in and manually change the business hours whenever there is a change between DST and standard time.
 - **Business Hours Reporting**—When business hours reporting is enabled, IC will use the time range you configure for the report and the business hours configuration information to determine what time range to use to generate the report (and to display on the report headers). If you enabled or changed the business hour configuration IC will use the settings that were in effect at the time the data for the selected time range was collected. However, if a change in settings occurred during the selected time range, IC will use the latest settings.
 - **Business Hours Reporting**—If you run a report that includes a time series plot using a time range, the timestamp for the plot will always be displayed as ending at 00:00 (midnight). For example, if your business hours set from 5:00PM to 2:00AM on 1-1-2011 and you run the business report on 2-1-2011, the time series will show that it is reporting data starting 5:00PM on 1-1-2011 to 12:00AM on 2-1-2011.
 - **Class Utilization report**—If you set the Network Group to a view that contains a sub-group, that sub-group will be excluded from the report. If you want to include sub-group information in the report, select the corresponding group or sub-group rather than selecting a view.
 - **Consolidated Reports**—Occasionally, when you select a consolidated report, the scheduling options are not enabled. To work around this issue, select a different consolidated report and then re-select the report you want to schedule; the scheduling options should now be enabled.
 - **Data Display**—In IC, the number of data points that get plotted depends on the data granularity and data retention settings you define on the Data Collector. When you change the granularity, you may initially notice some inconsistencies in the way data is plotted while DC adjusts.
 - **Data Display**—On a PacketShaper appliance, the average rate for a class or link may be greater than the partition size or link size if Shaping is OFF. In this case, IC reports will also show the partition utilization / link utilization percentage as greater than 100%.
 - **Data Display**—In IC reports, the Partition Burst Limit field corresponds to the data collected in the partition-burst-limit-bps ME variable. This variable represents the configured burst limit and, if the partition on the PacketShaper is not burstable, it is the same as the partition size. Therefore, IC cannot distinguish between burstable traffic and non-burstable traffic. Hence, if the PacketShaper appliance reporting the data does not have a burst partition, the Partition size and the Partition Burst Limit in IC reports will be identical.
 - **Data Display**—In some cases, an IC report may show average rate values in excess of peak rate values. In addition, the peak rate displayed on some site-based reports will not correlate with the peak rate of the parent class reported by the associated site on the PacketShaper. Both of these issues are due to an error in the data reported by the PacketShaper appliances.
 - **Data Display**—IC reports do not visually show breaks in data due to a device outage. For example, if a PacketShaper goes down for a day, DataCollector will not be able to collect any data for that time. However, any reports that span the time range of the outage will simply skip the outage time. You will notice that the time is missing on the report's x-axis, but you will not see a break in the lines on the graphs.

-
- **Data Display**—If the sites you define in IC comprise multiple PacketShapers, the Peak Rate that is displayed on graphs and charts represents the peak rate across all Shapers in the Site, not the peak per shaper. Therefore, the value on the Shaper with the highest peak rate in the site will be used as the Peak Rate for the site.
 - **Data Display**—Depending on the amount of data and the number of columns contained in a given report, some of the report columns may be truncated on the report display.
 - **Data Display**—If you move a traffic class on a PacketShaper appliance that DC has already started collecting from, it will no longer collect or report data for that traffic class. Additionally, if the traffic class is a site-based class on which you have created an IC site, DC will no longer report data for the associated site.
 - **FDR Display**—When you drill down to FDR data, the time range you configured for the report no longer applies. This is because reports only display raw FDR data. Therefore, the only data that is available is the data that is stored in the raw table. This is determined by the retention period that is configured for raw data.
 - **Link Utilization report**—When viewed in Microsoft Word (.doc) format, this report does not display the entire report initially. To fix this problem, simply reset the margins until the entire report is visible.
 - **Link Utilization Report**—The trending graph is calculated using all data that has been collected during the selected historical period regardless of whether business hours restrictions are in effect.
 - **ME Reports**—The data points and tick marks fall out of time selection, range and period when data is purged. This happens only in the first hour of data collection, where the data may get collected for the previous hour. For example, if the system starts data collection between 10 am to 11am, the user may see the data for last 15-20 minutes from 9am to 10 am. But, when the data is collected for another hour, for example, 11am onwards, this does not happen again, the user can view the data according to the selected time span.
 - **Scheduled Reports**—When a scheduled report runs, the report is not automatically displayed on the *Archived Reports* list. You must refresh the list in order to see the report entry.
 - **Site Activity Report**—The Normalized Delay column on the Site Activity Report is calculated using the network-delay-msec and trans-bytes ME variables ($\text{Normalized Delay (ms)} = 1000 * (\text{network-delay-msec} / \text{trans-bytes})$). However, on the PacketShaper reports Normalized Delay is derived from the normalized-network-delay-avg ME variable. Therefore, you will not see a correlation between the Normalized Delay values on the PacketShaper and IntelligenceCenter. For more information on the ME variables used by IC, refer to the **Reference > Measurement Variables** section of the online help.
 - **Time Zone**—If a report does not show the data you expect for a given time frame, daylight saving time (DST) may be affecting the results. If a device is set to automatically adjust its time zone for DST, the actual time that is being used on the device may differ from what you might expect given its configured time zone setting. For example, although a device may have a time zone setting of -8:00GMT, it may actually have its time set to match the -7:00GMT time zone if DST is in effect and the device is set to automatically adjust for DST.
 - **Time Zone**—The dates displayed on IC reports are based on the time zone configured on the client from where the report was run, not the time zone of the IC server. For example, if the IC server is configured with time zone -0800 and the client is configured with -0500, the reports the client generates will correspond to -0500. To work around this, you can set the client time zone to the same time zone as the IC server or you can select the Other time range when running the report and set it to the time range that aligns with the IC time zone.
 - **Top Sites and Top Applications by Site**—When you run the Top Sites report or the Top Applications by Site report using the Last Hour time range within the first 15 minutes of an hour (for example, if you run the report at 8:08), an exception may occur when plotting the pie-chart graph. This is because

the data has not yet finished rolling up into the database tables. As a workaround, re-run the reports after 15 minutes of the hour have elapsed (for example, try running the report again at 8:20).

- **Top Applications Report**—Sometimes the Top Application L3 report from Top listener report displays only one point. This happens due to an overlap of data points on the graph for some applications because the query returns the same result for them.
- **Top Services Report**—In Datacollector, all unknown services that do not have any collaboration (peers) are placed in the “Unknown-1” bucket. Thus, the drill down report from “Unknown-1” does not yield any peer data.
- **Top Immediate Children Report**—When you run the top immediate children report, the talker and listener information is provided for only the leaf node traffic classes. If the node selected in the report contains a non-leaf node, the talker and listener links does not display any record.
- **Traffic Class Compression Report**—If you run this report using a time range that is not using GMT 0000, the data on the time series chart in the fourth level of the report will not match the data on previous level. This is because if you have a GMT offset, the tables and charts on the previous levels of the report fetch data from both the “day” table and the “8 hour” table. However, the time series report can only fetch data from one of the tables.
- **View Reports**—When IC emails a PDF report using the email notification feature, the drill-down links in the report will not work because drill down requires access to IC and DC, which is not available from the email application.
- **View Reports**—In PowerPoint versions of reports, traffic class names look like hyperlinks, but they are not. Similarly, Microsoft Word versions of the reports do not display the icons for top talkers, listeners, or applications.

Supplementary Information

Installation/Upgrade Notes

- **Data Display**—You may notice a mismatch in the number of bytes, packets, and flows reported on Top N reports between report levels (i.e. between level 1 and level 2 and between level 2 and level 3) in version 3.3 of the product. In previous versions of the software, all top talkers, listeners, host pairs, services, VLAN IDs and DCSP values were rolled up from the raw data into the database hourly, daily, monthly, and yearly tables. However, in order to enhance reporting performance, IC 3.3 only stores the top 500 values (based on number of bytes) when it rolls data up. You may see this mismatch in all host-based Top N reports, such as the Top Listeners, Top Talkers, Top Sites, and Top Host Pairs.
- **Data Display**—For performance reasons, DC only rolls the top 500 services up into the hourly, daily, monthly, and yearly tables. And, for each of the 500 services, DC only keeps information about the top 100 host pairs that correspond to the service. However, for consistency, all other services are reported as “Others.” In addition, if you are upgrading from a previous release in which the full service and host pair information was available, this data will no longer be displayed in the reports you run because the logic to display it has been removed. Instead, it will display “Others” for host pairs beyond the limit.
- **Data Display**—In some rare instances, there is a delay in receiving flow data from PacketShaper appliances. If this delay is longer than 15 minutes, the data no longer makes it in to the raw data table and it will therefore not be rolled up into subsequent tables (and therefore it will not show up on reports).
- **Log File**—You may notice the following error in the DC installation log:

```
ERROR [org.apache.catalina.startup.Embedded] Cannot find specified temporary  
folder at C:\Windows\system32\config\systemprofile\AppData\Local\Temp\
```

This is normal and does not impact system performance.
- **PostgreSQL**—The PostgreSQL password will be set for you automatically.
- **Port Numbers**—Make sure that the port numbers you plan to use for HTTP and HTTPS access to IC are available before you install the software. If you have another service installed that uses the default port numbers for HTTP and HTTPS (IIS for example) and you install IC on the system using the default port numbers, you will not be able to launch IC. This may be true even if the other service is disabled.
- **Scheduled Reports**—If you upgrade from a previous IC 3.1.x version to IC 3.3.1.1, any reports that were previously scheduled will continue to run (unless the report is no longer available in the new release. In this case, you must manually remove the associated scheduled task). Note, however, that scheduled reports will not run until after the database migration process completes. Any reports that are scheduled to run during the migration process will be skipped. In addition, if you run into performance issues with IC, you should examine the list of scheduled tasks (**Manage > Scheduled Tasks**) and delete or deactivate any reports that you are not using.

IC-DC Interaction

You can manage DC from one IC server only. Do not add the same DC to the IC topology on more than one IC server. Doing so will break the reporting mechanisms. This limitation also prevents you from moving your IC server to a different machine or reinstalling IC. In order to move your server, you must back up your DC and IC databases using the IC and DC backup scripts (or using PostgreSQL backup tools). You can then reinstall the software and use the restore the databases to reinstate interaction between DC and the IC server. Keep in mind that IC and DC must be running the same version. That is, you cannot run a 3.3 IC server with a 3.2 DC.

Device Management

If you add a PacketShaper device to IC and later change the administration credentials for the appliance (for example by adding or changing the touch password), you must go back into IC and change the credentials on the device entry. To do this, select **Configure > Network > Devices** and select the appliance you want to modify.

PostgreSQL

- Both IC and DC use the PostgreSQL Database Server application to create their databases. For each connection to the Postgres database, a *Posgres.exe* process is opened. Therefore, you may see a large number of *Postgres.exe* processes (typically 20-30) show up on the Windows Task Manager **Processes** tab. This is normal and does not impact system performance.
- The PostgreSQL Database Server 8.4 service must be running on the server(s) where IC and DC are installed in order for IC and DC to function. If the PostgreSQL Database Server 8.4 service stops running, you must restart it as well as any IC and/or DC services that are running on the server. You must start PostgreSQL Database Server 8.4 before you can start the IC and DC services.

Authentication for ME Queries

Whenever DC sends an XML query to a PacketShaper appliance or to a PolicyCenter server that is managing appliances to collect measurement (ME) data, it must send three separate authentication login requests (one each for class, link, and partition data). Additionally, depending on the number of traffic classes running on each appliance, DC may need to send multiple XML queries in order to retrieve all of the ME data it needs. Therefore, you may see a large number of DC authentication requests in the PolicyCenter audit log. This is normal and does not impact system performance.

Collection of Missed Data

If DC loses connectivity to one of its ME data sources there will be gaps in the data that gets reported for that data source. When connectivity is reestablished, DC will automatically go back and collect the data that was missed during the outage, up to a maximum of 24 hours. However, depending on the length of the outage and the amount of data that was missed, it may take a while for DC to catch up on the data. You may notice discrepancies in reported data while DC is catching up on missed data. In addition, if the loss of connectivity was due to a Shaper outage, there may be some missed data that is not recoverable.

Reporting

- **Data Display**—DC polls its ME data sources for data according to the configured query interval. By default, DC polls ME data sources every 15 minutes. When you first start data collection, DC skips the first query interval to ensure that it receives a complete interval of data. For example, if you started DC at 10:32, it would set its first collection point to the start of the next interval, which begins at 10:45 (assuming a default query interval). It would then collect its first set of data when the interval completes at 11:00. If you then tried to generate a report with a time span that includes the period between 10:32 and 10:45 there would not be any data. These gaps in reporting should only be noticeable when data collection first starts. In addition, DC does not roll ME data up into the hourly table until 17 minutes after the hour, so the complete data set will not show up in ME reports for the Last Hour reporting time range until 20 minutes after the hour (that is, all Application, Device, and Site reports).
- **Data Display**—When comparing data across reports, consider that some reports are based on FDR data, some on ME data. For best correlation between reports, you should collect both ME and FDR data from the same set of devices. Keep in mind, however, that because FDR data is sent over UDP, delivery is not guaranteed and there may therefore be some slight differences between the ME and FDR reports. In addition, because many reports allow you to drill down on data, adding data filters

at each level of the report, you must make sure you are actually looking at the same “slice” of the data when comparing values across reports.

- **Data Display**—For usability reasons, pie charts are only displayed on Top N reports when N is 25 or less.
- **Log Files**—Right after installation, you may notice “Object not found” errors in your DC log files. These messages indicate that DC cannot find ME data that would have been collected in the hour before you installed IntelligenceCenter (and which, therefore, was not collected by DC). These messages are normal and will cease after a couple of hours of data collection.
- **Log Files**—You may notice the following errors in your DC log files:
`CollectorRollupMgr::Enqueue:Table name:c2010 is not valid!`
This is normal and does not impact system performance.
- **Log Files**—When a report runs, a TupleDesc reference leak warning message is captured in the postgres.log file. This is normal and does not impact system performance.
- **Scheduled Reports**—If you schedule a report using an existing schedule definition and subsequently modify or delete the schedule definition, any reports that were scheduled to run using that definition will continue to run unless you deactivate or delete them.
- **Scheduled Reports**—If you schedule a report or configure a portlet to report on a specific application, network view, or group/sub-group and later delete the application, view, or group/sub-group, the scheduled report or portlet configuration does not change automatically. You will need to manually reconfigure the portlet or scheduled report.
- **View Reports**—When a report runs and is added to the *Archived Reports* list, the first page of the report is static based on data that existed in the queried database tables at the time that the report was run. Drilling down from this static top-level view results in a live query to the database. If the data in the database tables that are queried has changed since the report was initially run, the top-level view of the report may no longer match the drill-down view of the report. For example, because IC provides the flexibility to allow you to report on any time range, you may generate reports for which data has not yet been completely rolled up into the database table that IC queries. In this case, drilling-down later will allow you to see additional data that was not yet available (and therefore not displayed on the top-level view) when the report was archived. Similarly, as time passes and data is rolled up, the IC database may no longer contain data that existed when a report was archived. This is especially true of old reports that have a short time span. Therefore, when you drill down from the top-level view, the resulting views may no longer contain any data.
- **View Reports**—When you choose to view a PDF version of a report in the report archive, only the top-level of the report will display in PDF. Drilling down to subsequent levels of the report can only be achieved from the HTML format, because drill downs require a live database query.

Application Performance Portlet Calculation

The Application Performance portlet gives information about application performance based on three network quality parameters: loss, latency, and availability. If there is not any data for any of the ME variables that are used to determine each of these parameters, the portlet will not display any data. However, if there is data for some of the parameters and no data for other parameters, IC will assume the full amount for the parameters for which it has no data in calculating whether the application conforms to the defined thresholds. Refer to the online help or the user guide for details about the Application Performance calculation.

Application Discovery

Many reports that you can run in IC detail the performance of applications on your network. Applications map to PacketShaper traffic classes, however they are not the same as traffic classes. Application traffic will not be displayed in an IC report or portlet if there is not an application defined in IC to represent it. By

default, IC includes applications that match many standard network traffic classes. The matching rules for each application, as related to the PacketShaper class tree, use the `/Inbound/*/<class_name>` and `/Outbound/*/<class_name>` format. With these matching rules, any standard traffic class will match the application definition as long as it is a leaf class. For example, the class `/Inbound/HTTP` and `/Inbound/Paris/HTTP` would both match the default HTTP application definition. However, until the traffic class discovery process completes, applications may show as inactive even though you have corresponding traffic running on your network. The amount of time required for class discovery depends on the number of data sources and traffic classes. During this discovery process, the reports you generate may not accurately display information about all applications and traffic classes running on your network.

IntelligenceCenter Interface

- **Data Sources**—IC allows you to define any level of the network topology or any network view as a data source. Any PacketShaper appliances and/or NetFlow-5 capable network devices (such as routers) in the network group, sub-group, or view are then automatically added as data sources and use the data collection settings you define for the group, sub-group, or view (metric or flow data). Because you can add views or groups at any level of the hierarchy, an individual device may be added as a data source multiple times, each with different data collection settings (for example, if you add a group and a sub-group as separate data source entries; if you add a sub-group and an individual device within the sub-group; or if you add a sub-group and a view that contain overlapping devices). The settings you define at the child level will always take precedence over the settings defined at higher levels of the hierarchy. For example, suppose you want to collect metric data from all PacketShaper appliances in a sub-group and you want to collect flow data from a single PacketShaper in the sub-group. In this case, you could add the sub-group as a data source and configure it to collect metric data. You could then add the individual PacketShaper from which you want to collect flow as a data source and configure it to collect metric and flow data.
- **Devices**—When you select a PacketShaper in the IC network, the **General** tab displays **Link Size** statistics even when shaping is turned off on the PacketShaper. Note that if shaping is off, the link size is not enforced.
- **IC Interface Panes**—When you select an object in the left pane of the IC user interface, IC displays the object properties in the right pane. Selecting a different tab in the left pane does not change the display in the right pane. For example, if you select a device entry on the **Topology** tab, the properties for the device are displayed in the right pane. If you then select the **View** tab in the left pane, the device properties continue to display in the right pane until you select another object in the left pane.
- **Required Fields**—In IC, fields that require input of a specific format will be highlighted in red until the proper string is entered. Sometimes fields will retain the red highlight even when a string of the correct format has been entered.

Single Sign-On for Radius and TACACS+ Users

Users can only use the single sign-on feature to access appliances or applications on which they have a local user name and password configured. Users who access the appliance or software via a RADIUS or TACACS+ account cannot use the single sign-on feature. Additionally, you will not be able to import devices from a PolicyCenter server if PolicyCenter authenticates those devices using RADIUS or TACACS+.

Port Numbers

The following table shows the port numbers used by IC and DC. For more details on how to prevent firewall problems with IC and DC, refer to Knowledge Base article FAQ973 (<https://kb.bluecoat.com/index?page=content&id=FAQ973>).

IP Address Requirements

Component	Port Number/range
IntelligenceCenter GUI	80 (insecure) or 443 (secure) by default. You can set these values to any available port number during the IC installation.
DataCollector	This is the port that IC and DC use to communicate. It is set to 8543 and cannot be changed.
Flow data collection	This is the port that DC listens on for FDR (Packeteer-2 and/or NetFlow-5) data. The default is 9800; the acceptable range is 1024-65535. All devices that are emitting FDR to DC must use the same UDP port number. You set the port number on which DC listens for FDR data on the Configuration tab for the DataCollector device entry (Configure > Network > Devices).
Metric data collection	DC collects ME data from PacketShaper appliances using XML over HTTP or HTTPS. If the Shaper is using HTTP, the port number must be 80. If the Shaper is using HTTPS, the default port is 443 and the range is 1-65535. You set the port number that IC and DC use to communicate with each Shaper on the General tab for the specific device (Configure > Network > Devices).
Property information synchronization between IC and Shaper devices	Every time IC requests property information from a PacketShaper device—either upon initial device deployment, user-initiated device refresh from the IC GUI, or when the IC does its regular device property synchronization—IC issues a passive FTP (unsecured) get request over port 21. If your IC server and the Shaper devices are separated by a firewall, you will need to open FTP port 21 (unsecured) on the firewall in order for IC to interact with and report on the device.
Local Postgres Database access	The Postgres database uses TCP ports 8778 and 5432 on the localhost address.

If you install IC and DC on separate systems, you must use static IP addresses on each system. This is because communication between IC and DC is established when you add DC to the IC topology based on the IP addresses on each system at that time. If the IP address on either system changes after you deploy DC, you will have to remove DC from the topology and re-add it in order to reestablish communication. Keep in mind that if you remove DC from the topology, you will lose all configuration settings.

In addition, if you change the DC address after deployment, any PacketShaper appliances or other network devices that you have configured to emit FDR to it will also have to be reconfigured.