

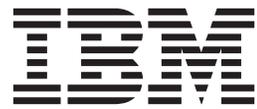
IBM Watson Explorer
Version 9.0.0.5

Installation and Integration Guide



IBM Watson Explorer
Version 9.0.0.5

Installation and Integration Guide



Note

Before using this information and the product that it supports, read the information in “Notices and trademarks” on page 85.

Contents

Installing and Integrating IBM Watson

Explorer. 1

Overview of Watson Explorer and Related Modules	1
Requirements	1
Server Installation Requirements.	1
Satisfying Installation Requirements for SUSE Linux.	3
Server Hardware Recommendations	4
Client Requirements	4
Installing Watson Explorer Product Modules.	6
Overview of Watson Explorer Installation Models	6
Pre-Install Configuration and Data Collection	6
Selecting and Preparing Installation Media	12
Overview of the Installation Process	15
Installing Watson Explorer Graphically	15
Installing Watson Explorer from the Command-Line	23
Performing Post-Installation Configuration	35
Troubleshooting an Installation	45
Getting Help	46
Using Response Files for Silent Installation	47
Locating Response Files	47
Modifying Response Files	47
Installing Using a Response File	48
Modifying Watson Explorer Installations.	48
Reinstalling Watson Explorer	49
Uninstalling Watson Explorer and Installation Manager	50
Uninstalling Watson Explorer Modules	50
Uninstalling Installation Manager	50
Uninstalling Watson Explorer Modules using a response file	51
Upgrading from Previous Releases	52

Upgrading from Watson Explorer 9.0 to Fix Packs	52
Upgrading from Watson Explorer 8.2 Installations	54
Upgrading from Velocity to Watson Explorer Engine	55
Upgrading from IO Pro to Watson Explorer Results Module	56
System Administration	56
Identifying Installed Watson Explorer Software Versions	57
Environment Variable Settings for Watson Explorer Engine	58
System and Application Configuration for Watson Explorer Engine	60
Configuring Watson Explorer Engine's Embedded Web Server	61
Command-Line Watson Explorer Engine Administration Commands	68
Integrating Watson Explorer Into the System Startup and Shutdown Process	75
Integrating Watson Explorer with Anti-Virus Software	80
Integrating Watson Explorer Engine with System Cleanup Scripts	80
Integrating Watson Explorer Engine with System Updates	81
Backing Up Watson Explorer Engine	81

Notices and trademarks 85

Providing comments on the documentation 89

Installing and Integrating IBM Watson Explorer

This portion of the documentation identifies the requirements for successfully installing and running Watson Explorer, explains how to install one or more of the modules that make up Watson Explorer, and discusses how to integrate Watson Explorer modules into your system administration environment.

Overview of Watson Explorer and Related Modules

IBM® Watson™ Explorer is available in two versions, a Standard Edition (SE) and an Enterprise Edition (EE). These version differ only in terms of how they are licensed and the pricing for those capabilities. Both of these packages contain exactly the same software.

Both of these Watson Explorer offerings are comprised of the following installable modules:

- **Application Builder:** Simplifies developing 360-degree information applications that deliver targeted, relevant data to users based on their roles or interests within the enterprise or a specific organization
- **BigIndex:** Provides a Java API for creating, configuring, and deploying information navigation applications in enterprise environments. The BigIndex API includes sophisticated functions that simplify or eliminate traditionally manual enterprise application administration tasks such as load balancing, mirroring, and replication.
- **Watson Explorer Engine:** Provides a web-based development and deployment environment for enterprise information navigation applications
- **Knowledge Center:** Provides online help for Watson Explorer Application Builder, Watson Explorer Engine, and Results Module. By default, online help within these applications contacts a centralized IBM server to deliver online help. Installing and starting the Knowledge Center locally provides online help for Watson Explorer in "dark" environments where development systems may not have access to the Internet, and in environments where Internet access may be slow or inconsistent.
- **Results Module:** Enables local administrators of enterprise information navigation applications to easily extend those applications with business-specific highlighting of selected query results and by incorporating domain-specific knowledge such as taxonomies
- **ZooKeeper:** A distributed coordination service and associated data repository for storing configuration information about distributed applications.

Requirements

The following sections explain the requirements for installing any Watson Explorer module(s), provide recommendations for minimum hardware configurations for development and deployed systems, and discuss client requirements for accessing the administrative interfaces for any Watson Explorer module(s) and applications that have been created using those modules.

Server Installation Requirements

Systems must satisfy the following requirements in order to install and successfully execute Watson Explorer and related product modules:

- **Supported 64-bit hardware:** A 64-bit (AMD64 or Intel 64) x86 system with a 2 GHz or better processor. A minimum of 4GB of memory is recommended.

Important: When running any Watson Explorer module on a system with less than 16 GB of memory, that system should be configured to use a swap file or partition in order to be able to support virtual memory. Follow your operating system's recommendations for the size of a swap file or partition, and see your operating system's documentation for information about how to add swap space.

- **Supported 64-bit operating system:**
 - 64-bit Red Hat Enterprise Linux (RHEL) 5.3 Server or Advanced Platform distribution or later
 - RHEL 6.0 Server or later
 - Equivalent 64-bit RHEL-based distributions
 - SUSE Enterprise Linux Server 10 or greater
 - 64-bit version of Microsoft Windows Server 2008 Standard, Datacenter, or Enterprise editions
 - Microsoft Windows Server 2012 Standard or Datacenter editions

Note: Watson Explorer is expected to work correctly on Windows Server 2012 R2, but neither that version of Windows Server nor Microsoft IIS 8.5 were tested.

- **Installation Manager 1.6.3.1 or greater:** The Watson Explorer installation process uses IBM's Installation Manager to install Watson Explorer and track installed versions. You do not need to separately download or install Installation Manager. A bundled version of Installation Manager is included on the DVD, and is used and installed during the Watson Explorer installation process if Installation Manager is not already installed on your system.

Note: When installing any Watson Explorer module on a Linux system using the `launchpad.sh` or `launchpad-console.sh` scripts, you may see a message which reports that the library `canberra-gtk-module.so` can not be found. The absence of this library does not cause any problems when installing or subsequently using the version of Installation Manager that is bundled into the ISO image or zip archive.

- **Version 10 or greater of the Microsoft Visual C Runtime:** To use IBM Text Analytics on a Windows server you must have Version 10 or greater of the Visual C Runtime installed. The Runtime is offered as a part of the redistributable Visual C++ binary from Microsoft and can be downloaded here. You only need to install this Runtime if you plan to use IBM Text Analytics on a Windows server.
- **Internet connection:** Installing Watson Explorer from downloadable media requires that the system on which you are installing any Watson Explorer module(s) can access the Internet.
- **Web Browser:** Installing Watson Explorer in graphical mode requires that a supported web browser is installed as the default browser for the graphical desktop environment that is used on your system. The default browser on Microsoft Windows systems is Internet Explorer. The default browser on Linux systems is typically Firefox. See "Client Requirements" on page 4 for information about supported browsers and browser versions.
- **Web server:** Watson Explorer Engine includes an embedded web server, but is most commonly used in enterprise environments with system web servers such as Apache or Microsoft Internet Information Services (IIS).

When integrating with an enterprise web server, while Watson Explorer Engine has been shown to run on a wide variety of web server software, IBM recommends using one of the following:

- Apache 2.x
- Microsoft IIS 6
- Microsoft IIS 7 or 7.5
- Microsoft IIS 8

Note: Watson Explorer Engine is expected to work correctly on Microsoft IIS 8.5, but neither that version of IIS nor Windows Server 2012 R2 were tested.

Note: The Watson Explorer Results Module and Watson Explorer Application Builder execute in the context of an application server, include a bundled version of IBM WebSphere Liberty Profile, and therefore do not require integration with a system web server.

Satisfying Installation Requirements for SUSE Linux

As part of the installation process for Watson Explorer and related modules, you will be prompted for the Linux user and group who will own the files and directories in the final installation. The user that you specify must be a member of the group that you specify. In addition, the user that you specify must be associated with a functional Linux shell if that user is to be able to start the services that are associated with various Watson Explorer and related modules.

You may need to make two modifications to your system in order to successfully install and run Watson Explorer and related modules on your system:

- **Verify Correct User and Group Membership for the user who will own your Watson Explorer installation:**

On some versions of SUSE Linux, the Apache web server is installed as a user (**wwwrun**) who is not a default member of the group that the Apache web server runs as (**www**). If you are installing Watson Explorer Engine in administrator mode on a SUSE system, you must correct this before starting the actual installation process.

To determine whether this user is a member of the required group, execute the following command:

```
grep www /etc/group
```

If you see output like the following, your system's user and group configuration is correct:

```
www:x:8:wwwrun
```

If the **wwwrun** user is not displayed on this line, you will need to add the **wwwrun** to the **www** group in order to proceed with the installation.

How you add a user to a group on a Linux system depends on the authentication model that you are using on that system. If you are using local password and group authentication, you can add the **wwwrun** user to the **www** group by executing the following command as the root user or via the **sudo** command:

```
groupmod -A wwwrun www
```

If your SUSE systems use a different authentication model, you will have to perform equivalent commands to ensure that the user as whom your web server is running is a member of the group that the web server is running as.

- **Verify the shell used by the user who will own your Watson Explorer installation:**

In order to be able to start services, a user must have a login shell that enables the execution of other processes. Though the **wwwrun** user on SUSE systems does not have a password that would enable that user to log in, that user is assigned **/bin/false** as their default shell to further ensure system security.

In order to enable the **wwwrun** user to execute processes, you must change that user's default shell from **/bin/false** to a valid shell, such as **/bin/bash**. You can do this by executing the **chsh** (Change Shell) command as the root user or via the **sudo** command, as shown in the following example:

```
chsh -s /bin/bash wwwrun
```

After verifying correct user/group membership for the user that you will be install Watson Explorer and related modules as, and ensuring that the specified user has a valid login shell (even if they cannot actually log in), you are ready to proceed with installing Watson Explorer and related modules.

Server Hardware Recommendations

The previous section discussed the minimum hardware and operating system requirements for a server system on which you want to develop or deploy Watson Explorer applications. Simply satisfying these minimum requirements will not provide an optimal development or deployment experience.

IBM's suggested minimum hardware and operating system configurations for Watson Explorer development and deployment are the following:

- **Development system:**
 - Single Dual-Core 2Ghz 64-bit CPU
 - 4GB RAM
 - 250GB disk space
 - Supported 64-bit Linux or Microsoft Windows operating system
- **Deployable Watson Explorer Engine Server:**
 - Dual Quad-Core 2Ghz 64-bit CPU
 - 16GB RAM
 - 500GB disk space, RAID5
 - Supported 64-bit Linux or Microsoft Windows operating system

Note: The different modules of the IBM Watson Explorer product are designed so that they can be installed and used on physically distinct server systems. Installing all Watson Explorer modules on a single server is not a recommended deployment model.

Client Requirements

The only requirement for IBM Watson Explorer modules is a modern web browser.

Each module has slightly different client requirements. For more details about browser support and other support information, see the following table:

Table 1. Client requirements for Watson Explorer modules

Module	Supported browsers and other requirements
Watson Explorer Application Builder	<ul style="list-style-type: none"> • Firefox 17 or better • Microsoft Internet Explorer 8 or better <p>No issues have been reported when using other modern standards-based browsers such as Safari 3 or better, Opera 8 or better, and Google Chrome.</p>
Watson Explorer BigIndex	Not applicable. BigIndex provides only an API and associated libraries.
Watson Explorer Engine	<ul style="list-style-type: none"> • Firefox 3.6 or better • Microsoft Internet Explorer 6 or better <p>No issues have been reported when using other modern standards-based browsers such as Safari 3 or better, Opera 8 or better, and Google Chrome.</p> <p>Using the Watson Explorer Engine administration tool in a browser requires that browser support for JavaScript is enabled.</p> <p>The Watson Explorer Engine administration tool is not designed to be used through a reverse proxy, and is not supported in that environment.</p> <p>Important: Support for Internet Explorer versions 6 and 7, and Firefox versions lower than Firefox 18 will be discontinued in a future release of Watson Explorer Engine.</p>
Watson Explorer Knowledge Center	<ul style="list-style-type: none"> • Microsoft Internet Explorer 9 or better • Firefox 17 or better • Chrome 20 or better • Safari 5 or better <p>You can use earlier versions, but because Watson Explorer Knowledge Center is optimized for the versions that are listed, some functionality might not be available in lower versions.</p>
Watson Explorer Results Module	<ul style="list-style-type: none"> • Firefox 3.6 or better • Microsoft Internet Explorer 7 or better <p>No issues have been reported when using other modern standards-based browsers such as Safari 3 or better, Opera 8 or better, and Google Chrome.</p>

Installing Watson Explorer Product Modules

The next few sections provide overview information about the ways in which those modules can be installed, and then provide a series of sections that walk you through the pre-installation, installation, and post-installation phases of exploring Watson Explorer. This section concludes with troubleshooting information and information about how to obtain assistance from IBM product support if you cannot resolve any installation, configuration, or integration problems that you are experiencing.

Overview of Watson Explorer Installation Models

IBM recommends that you install any Watson Explorer or related module as the root user on Linux systems, or as a user with administrator privileges on Microsoft Windows systems. However, you can also install Watson Explorer as a user who does not have administrator privileges, which is known as installing in "non-admin" mode. In order to install any Watson Explorer module in non-admin mode, you must also install Installation Manager in non-admin mode, as part of the installation process.

Graphical or command-line installation: Whether installing Watson Explorer as an administrator (admin) or non-admin user, you can install any Watson Explorer or related software module by using Installation Manager in command-line mode (using the **launchpad-console.sh** script) or in graphical mode (using the **launchpad.sh** script). Installing Watson Explorer in graphical mode is explained in "Installing Watson Explorer Graphically" on page 15. Installing Watson Explorer in command-line mode is explained in "Installing Watson Explorer from the Command-Line" on page 23.

Automated, silent installation: Installing any Watson Explorer or related software module generates a **response file** for that module which you can subsequently use to install that module on another system non-interactively, without requiring responses to the prompts that are normally displayed during the installation process. Response files provide the responses to each step of the installation process that were supplied during the installation in which that response file was generated, regardless of whether that installation was done using the graphical or command-line version of Installation Manager. Because installation using response files requires no interaction after initiating the appropriate command, using a response file to install any Watson Explorer or related software module is known as "silent installation". See "Using Response Files for Silent Installation" on page 47 for detailed information about using response files for silent installation of Watson Explorer and related modules.

Pre-Install Configuration and Data Collection

Installing one or more Watson Explorer modules on different platforms requires that you collect some information about those systems and configure any associated system servers appropriately. The next sections discuss how to ensure that a system provides the physical resources and performance that are required for Watson Explorer development and deployment, and then discuss the pre-installation requirements that are unique to different supported platforms.

Verifying System Capabilities and Resources

Before installing any Watson Explorer module, you should consider examining the configuration and capabilities of the system on which it is installed to ensure that it provides a suitable hardware and software environment for running the Watson Explorer module(s) that you are installing. System capabilities are especially critical

for systems on which you are installing Watson Explorer Engine, because the enterprise information applications that you create with Watson Explorer Engine can require substantial network and system resources in order to crawl remote data, index that data, search large indices, return large result sets, and so on.

To help determine whether a given computer system satisfies Watson Explorer hardware and performance requirements, IBM includes an application named **micro-test** to measure the actual real-world performance of any server or test system on which you plan to run any Watson Explorer module. When executed, the **micro-test** utility performs one or more program execution, read-write, and system analysis tests and records summary data about the results and performance of those test. These results can then be used by customers, product support, and professional services personnel to determine whether a given server is suitable for deploying one or more Watson Explorer modules and satisfy your requirements.

The **micro-test** utility (**micro-test.exe** on Microsoft Windows systems) is installed as part of any Watson Explorer Engine installation, and is located in the **bin** sub-directory of the product's installation directory. Documentation for the **micro-test** utility is available as part of the Watson Explorer Knowledge Center.

IBM strongly recommends that you execute the **micro-test** utility on each host where you have installed any Watson Explorer module, and that you do so and archive the results of its tests before starting any Watson Explorer module for the first time. This will enable the **micro-test** utility to collect information about the default performance of each prospective host, and will also alert you to the presence of existing system software, such as anti-virus software, that you will need to configure so that it does not interfere with the module(s) that you have installed. IBM recommends that you execute the **Test All** batch test, in order to collect data about all of the aspects of system performance and configuration that the **micro-test** utility can collect.

Note: The **micro-test** utility is designed to review current system capabilities and performance, and should not be considered a general-purpose benchmarking utility.

Enabling Software Installation on Linux Systems

Note: The instructions in this section are only necessary to perform if you are going to install Installation Manager and Watson Explorer as an administrator, into system directories such as /opt. You do not need to disable SELinux if you are going to install both Installation Manager and Watson Explorer in non-admin mode, which installs them into your home directory.

If you are running a version of Linux that includes support for AppArmor or Security-Enhanced Linux (SELinux), please disable them before installing any Watson Explorer module. You can disable SELinux policy enforcement while installing Watson Explorer modules by executing the following command as the root user:

```
/usr/sbin/setenforce 0
```

Executing this command will put SELinux in permissive mode, where violations of SELinux policies are logged but not prevented, until the next time you boot your system or until you execute the following command, which will put your system back in enforcing mode:

```
/usr/sbin/setenforce 1
```

All Watson Explorer modules require access to resources that are typically managed by SELinux, and will therefore not function correctly on systems where SELinux is actively enforcing its default policies. Contact product support for information about creating SELinux policies that enable the Watson Explorer modules to coexist with active SELinux policy enforcement.

Using the Right Version of Installation Manager

If Installation Manager is installed on the system that you will install IBM Watson Explorer on, to ensure that the product installs correctly, you might need to upgrade to the version that Watson Explorer uses.

Watson Explorer uses IBM Installation Manager to provide the framework for installing and uninstalling one or more Watson Explorer modules. Installation Manager is widely used as an installation framework for many IBM products.

Each new release of Installation Manager adds capabilities that were not present in previous releases. Installation Manager is therefore sensitive to differences between the version of Installation Manager that is bundled with products as their installer, and versions of Installation Manager that have already been installed on a system as part of installing other products.

Important: IBM recommends against using Installation Manager 1.7 to install Watson Explorer.

Perform the following task to determine if you have Installation Manager installed and, if so, whether you need to install a later version:

1. Check your system to see if you have the IBMIM executable file:
 - For Linux systems: The default installation directory is `/opt/IBM/InstallationManager/eclipse/IBMIM`
 - For Microsoft Windows systems: The default installation directory is `C:\Program Files\Installation Manager\eclipse\IBMIM.exe`
2. If you do not have the executable file, Installation Manager is not installed on your system. You do not need to complete the remaining steps in this task, and can continue with other necessary preinstall configuration tasks. When you install Watson Explorer, the correct version of Installation Manager will be installed.
3. If you have the executable file, Installation Manager is installed on your system. Proceed with the next step in this task.
4. Run Installation Manager:
 - For Linux systems, with administrator privileges, enter **`opt/IBM/InstallationManager/eclipse/IBMIM`**
 - For Windows systems, from the `C:\Program Files\Installation Manager\eclipse` directory, run `IBMIM.exe`
5. From the **File** menu, select **Preferences**.
6. From the left pane of the Preferences window, select **Updates**, and ensure that **Search for Installation Manager updates** is not selected.

Important: As explained previously, specific product versions are bundled with specific versions of Installation Manager. If you update the version that is installed, it is possible that an updated version of Installation Manager might not work the same way as a previous version.

7. From the **Help** menu, select **About IBM Installation Manager**.

8. If you have Installation Manager, Version 1.6.3.1, you have the version of Installation Manager that is bundled with Watson Explorer. You do not need to complete the remaining steps in this task, and can continue with other necessary preinstall configuration tasks.
9. If you have Installation Manager, Version 1.7:
 - If Watson Explorer is not already installed on your system, you do not need to complete the remaining steps in this task, and can continue with other necessary preinstall configuration tasks. You can use Installation Manager 1.7 to install Watson Explorer only if it was never installed on your system.
 - If Watson Explorer is installed on your system, contact product support for information about installing a new version of Watson Explorer. For information about contacting support, see “Getting Help” on page 46.
10. If you have a version of Installation Manager that is earlier than 1.6.3.1, Download the 32- or 64-bit version of Installation Manager 1.6.3.1 from the IBM Support Portal and install that version on your system.
11. Add the Watson Explorer repository to the repositories that are used by Installation Manager.
 - a. Start Installation Manager on your system and select **File > Preferences**.
 - b. Click **Repositories** from the left-hand pane to display the list of repositories that are currently available to Installation Manager.
 - c. Click **Add Repository**.
 - d. In the Add Repository window that opens, browse to the directory in which you extracted the contents of the Watson Explorer archive.
 - e. Browse to the ide-offering subdirectory of the extracted Watson Explorer.
 - f. Select the `diskTag.inf` file in this directory and click **Open**.
 - g. Click **OK** to close the Add Repository window.
 - h. Click **OK** to save the updated preferences setting and close the Preferences window.Continue with other necessary preinstall configuration tasks.

Pre-Installation Tasks for Watson Explorer Engine

The next few sections explain information and configuration that is required to successfully install and configure Watson Explorer Engine on supported operating systems.

Collecting Configuration Information for Linux Installations: As part of the installation process for Watson Explorer, you will be prompted for the Linux user and group who will own the files and directories in the final installation. This is especially important for Watson Explorer Engine, which is either integrated into an existing web server installation or runs an embedded web server that is included with the product, but is a requirement of other Watson Explorer and related product modules because some of the files and directories that are used by Watson Explorer must be writable by both a specific user and a group which that user belongs to. The user that you specify during installation must also be a member of the group that you specify.

Note: If you are installing Watson Explorer or related software modules on a system that is running SUSE Linux, make sure that you have verified that the requirements discussed in “Satisfying Installation Requirements for SUSE Linux” on page 3 before beginning your Watson Explorer installation.

If you are installing Watson Explorer Engine in administrator mode on a Linux system, you will probably want to integrate it with a system web server such as

the Apache web server, and should collect the information discussed in “Collecting User/Group Information for Privileged Linux Installations” before beginning your Watson Explorer installation.

If you are installing Watson Explorer Engine in non-admin (un-privileged) mode, you should collect the information discussed in “Collecting User/Group Information for Non-Admin Linux Installations” before beginning your installation.

Collecting User/Group Information for Privileged Linux Installations:

As discussed in the **Web Server** section of “Server Installation Requirements” on page 1, Watson Explorer Engine has additional installation requirements than other Watson Explorer modules because it is commonly integrated with a system web server. Before installing Watson Explorer Engine on a Linux platform and integrating it with an existing web server, you will need to know the user and group under which the web server on your system runs. On modern Linux systems, these are typically the user and group **apache** or **www-data**. If you are unsure which of these users and groups is used on your system, you can find out by running the **ps** command at a shell prompt on your Linux system, as in the following example:

```
# ps -eo user,group,comm | egrep 'apache|httpd'

root      root      httpd
apache    apache    httpd
apache    apache    httpd
apache    apache    httpd
```

The three columns of this output identify the user, group, and name of every process that contains either the strings "apache" or "httpd". This output shows that, while the **root** user and group started the first **httpd** process, the others are running as the user and group **apache**, as expected. If you are running a web server other than the Apache web server (which is not recommended), you would change the processes that you are searching for (apache|httpd, in the previous example) to contain the name of your web server's executable.

Collecting User/Group Information for Non-Admin Linux Installations:

If you will be installing in non-admin mode on a Linux system, you will need to supply the name of a group of which the user who is doing the installation is a member. To display the list of all groups that this user is a member of, execute the **groups** command at a shell prompt on your Linux system, as in the following example:

```
# groups
users engdev documentation testing root
```

You will need to specify the name of one of these groups when prompted to do so during the non-admin mode installation process. You should generally specify a group that does not have any special system privileges.

Pre-Configuring Microsoft IIS on Windows Systems:

If you are installing Watson Explorer Engine in administrator mode, to install and use that module with Microsoft Internet Information System (IIS) 7 or higher web servers, specific configuration options must be enabled in those web servers.

Before you begin

If the following conditions are true, ensure that Microsoft Internet Information System (IIS) is installed:

- You will install Watson Explorer Engine in admin mode.
- You will not use the embedded web server that is included with Watson Explorer Engine.

Important: Regardless of the version of IIS that you use, you might need to manually configure some aspects of that virtual directory, as explained in “Post-Install Configuration on Windows” on page 36.

If you will install Watson Explorer Engine in non-admin mode on a Windows system, you do not need to configure IIS to enable integration of Watson Explorer Engine with IIS. Versions of Watson Explorer Engine that were installed in non-admin mode can be accessed only through the embedded web server that is included with Watson Explorer Engine.

Installing and configuring IIS is not required on systems on which you will install only the Results Module or Application Builder, because these Watson Explorer modules include their own WebSphere web servers.

About this task

When installing Watson Explorer Engine, the installer provides integrated support for creating virtual directories in IIS on systems that are running IIS 7 or higher. For IIS 7 or higher, the installer requires that certain IIS components are installed and activated in your IIS configuration. The following components are required:

- CGI
- ISAPI Extensions
- ISAPI Filters
- IIS Management Scripts and Tools

For IIS 6, see “Creating a Virtual Directory in an IIS 6 Web Server” on page 37.

To verify that the required components are installed and active, do the following on the system where IIS 7 or higher is installed and on which you will install Watson Explorer Engine:

Procedure

1. Open Server Manager:
 - Microsoft Windows Server 2008: Click **Start > Administrative Tools > Server Manager**. The Server Manager window displays.
 - Microsoft Windows Server 2012: From the desktop, hover your cursor in the upper- or lower- right corner of the screen to show the charms, click the **Start** charm, and then the **Server Manager** tile.
2. Select the web server:
 - Microsoft Windows Server 2008: In the left pane of the Server Manager window, expand **Roles** and select **Web Server (IIS)**.
 - Microsoft Windows Server 2012: In the left pane of the Server Manager window, select **IIS**.
3. Verify that the following role services are installed: **CGI, IIS Management Scripts and Tools, ISAPI Extensions, and ISAPI Filters**.
 - Microsoft Windows Server 2008: In the Web Server (IIS) pane, scroll all the way to the end of that pane to find the **Role Services** section. Role services are listed in alphabetical order by role, role group, and then role services:

- CGI, ISAPI Extensions, and ISAPI filters are located in **Web Server (IIS) > Web Server > Application Development**.
 - IIS Management Scripts is located in **Web Server (IIS) > Web Server > Management Tools**.
 - Microsoft Windows Server 2012: In the right pane of the Server Manager window, scroll all the way to the end of that pane to find the **ROLES AND FEATURES** section.
4. If some or all of the required role services are not installed:
 - Microsoft Windows Server 2008: From the **Role Services** section, click **Add Role Services**, select the required information in the initial pages of the wizard, and advance to the Select Role Services page.
 - Microsoft Windows Server 2012: From the **TASKS** menu of the **ROLES AND FEATURES** section, select **Add Roles and Features**, select the required information in the initial pages of the wizard, and advance to the Server Roles page.
 5. If any of the required role services are not installed, select them from the Roles list. The following table lists the required role services and their locations:

Role service	Location in the Roles list
CGI	Web Server (IIS) > Web Server > Application Development
ISAPI Extensions	Web Server (IIS) > Web Server > Application Development
ISAPI Filters	Web Server (IIS) > Web Server > Application Development
IIS Management Scripts and Tools	Web Server (IIS) > Web Server > Management Tools >

6. Advance to the confirmation page and click **Install**.

What to do next

If you had to install the CGI role service, restart IIS:

- Microsoft Windows Server 2008: In Server Manager, expand **Web Server (IIS)**, select **Internet Information Services (IIS) Manager**, and from the Connections pane, select the name of your web server. From the Actions pane, click **Restart**.
- Microsoft Windows Server 2012: From the desktop, hover your cursor in the upper- or lower- right corner of the screen to show the charms, click the **Start** charm, and then the **Internet Information Services (IIS) Manager** tile. From the Connections pane, select the name of your web server. From the Actions pane, click **Restart**.

At this point, you can proceed with the Watson Explorer Engine installation process.

Selecting and Preparing Installation Media

Watson Explorer and related modules can be installed from the following sources:

- physical DVD media that you ordered and received from IBM or a partner
- an installation archive file in zip format that you downloaded from IBM after authenticating using your Passport Advantage credentials

The next sections explain how to make those installation sources available on Linux or Microsoft Windows systems.

Accessing DVD Installation Media

This section explains how to make Watson Explorer DVD installation media accessible on Linux and Windows systems so that you can begin the installation process.

Mounting DVD Media on Linux Systems

If you received Watson Explorer from IBM on physical DVD media, inserting that DVD into a DVD drive on a Linux server may mount it automatically. If this is the case, the DVD will probably be mounted with a name of the form `/media/IDE_VERSION-EDITION`, where:

- *VERSION* is a text string that identifies the version of the Watson Explorer software that is on the DVD, such as **v9.0.0**
- *EDITION* is the edition of the Watson Explorer software that is on the DVD, usually either **SE** (Standard Edition) or **EE** (enterprise edition)

Listing that directory should show the files and directories that are located at the root of the ISO image or physical DVD. If this is the case, change directory to that directory on your system and proceed to the end of this topic for instructions on how to proceed with the installation process.

More commonly, auto-mounting DVDs is disabled on enterprise Linux systems. In this case, you must manually mount the device that contains the DVD on a directory on your system. To mount your DVD device on a Linux system, execute a command like the following as the root user or via **sudo**:

```
/bin/mount -t iso9660  
DVD-DEVICE MOUNT-DIR
```

In this example command, *DVD-DEVICE* is the name of the device node for the device in which you inserted the DVD, and *MOUNT-DIR* is the name of a directory on your system where you want to mount the DVD. The most common device node for a DVD drive is `/dev/sr0`, though this may differ if your system contains multiple CD or DVD drives. The typical directory on which new filesystems are mounted manually is the directory `/mnt`.

If this command executed correctly, listing the `/mnt` directory should show the files and directories that are located at the root of the physical DVD. If this is the case, change directory to the `/mnt` directory on your system.

Proceed to the end of this topic for instructions on how to proceed with the installation process.

Using Local Installation Media on Windows Platforms

On Microsoft Windows systems, inserting a DVD into a DVD drive will normally make the contents of that DVD available through the drive letter that is associated with the DVD drive. Most Windows servers are configured so that they do not automatically execute applications from removable media.

Proceeding with the Installation Process

To proceed with the installation process:

- If you are performing a graphical installation of Watson Explorer or related modules, proceed to “Installing Watson Explorer Graphically” on page 15
- If you are performing a command-line installation of Watson Explorer or related modules, proceed to “Installing Watson Explorer from the Command-Line” on page 23

Extracting the Contents of a Downloaded Installation Archive

The Watson Explorer installation media can also be downloaded as an archive file in zip format that contains the same hierarchy of files that is found on physical DVD media, and is therefore referred to as an **installation archive**. This section explains how to extract the contents of a downloaded Watson Explorer installation archive on Linux and Windows systems so that you can begin the installation process.

Expanding Zip Archives on Linux Systems

If you have downloaded a Watson Explorer installation archive for use in installing Watson Explorer or related modules on a Linux system, do the following to prepare for the actual install process:

1. Change directory to the directory where you downloaded the Watson Explorer installation archive and execute the following command from a terminal window or console:

```
unzip IMAGE-FILE-NAME
```

Replace *IMAGE-FILE-NAME* with the name of the installation archive file that you downloaded. Executing this command creates a directory with the name of the release contained in the archive, such as `IDE_v9.0.0-SE` or `IDE_v9.0.0-EE`, in your working directory.

Note: If the **unzip** command is not installed on the system where you are installing Watson Explorer or related modules, you can install it by executing the following command as the root user or via the **sudo** command:

```
yum install unzip.x86_64
```

2. Change directory to the directory that was created in the previous step.

Proceed to the end of this topic for instructions on how to proceed with the installation process.

Expanding Zip Archives on Windows Systems

If you have downloaded a Watson Explorer installation archive for use in installing Watson Explorer or related modules on a Windows system, you can use Windows built-in support for archive files in zip format to extract the contents of that archive so that the extracted files and directories can be used for installation. To extract the contents of a zip archive file on Microsoft Windows systems:

1. Right-click on the downloaded archive file and select **Extract All...** from the context-sensitive menu that displays. The Extract Compressed (Zipped) Folders dialog displays.
2. Click **Browse** to navigate to the location where you want the files and directories that are contained in the archive to be extracted.
3. Click **Extract** to begin the extraction process. A progress dialog displays as the contents of the archive are extracted, and closed once the extraction process is complete.

Proceed to the end of this topic for instructions on how to proceed with the installation process.

Proceeding with the Installation Process

To proceed with the installation process:

- If you are performing a graphical installation of Watson Explorer or related modules, proceed to “Installing Watson Explorer Graphically”
- If you are performing a command-line installation of Watson Explorer or related modules, proceed to “Installing Watson Explorer from the Command-Line” on page 23

Overview of the Installation Process

Installation Manager's command-line and graphical modes share the same basic structure, consisting of the following basic steps:

1. **Select installation mode** - specify whether you are installing Watson Explorer and related modules in Administrator or User (non-admin) mode
2. **Install Packages** - specify that you want to install Watson Explorer and related modules
3. **Licenses** - enables you to review the licenses under which the applications and modules that you are installing are released
4. **Package and Shared Directory Location** - enables you to specify the location of a directory that will be used by Installation Manager to download artifacts and store information about the software packages that it has installed
5. **Features** - enables you to select the programs and modules that you want to install
6. **Custom panels** - displays custom panels that enable you to provide additional configuration information during or used by the installation process. The panels that you see depend on the features that you selected for installation, the operating system on which you are installing Watson Explorer, and whether you are installing Watson Explorer in privileged or non-admin mode.
7. **Summary** - displays summary information about the applications and modules that you have selected for installation, and any associated custom information that you have specified, before actually beginning the installation process

Installing Watson Explorer Graphically

After enabling local access to the installation media, you can install any Watson Explorer or related product module in graphical mode by following the instructions that begin in this section.

If you are installing from a DVD, inserting that DVD into a system's DVD drive may attempt to initiate the installation process, depending on how your system is configured to respond to media insertion. As explained in “Overview of Watson Explorer Installation Models” on page 6, you can install Watson Explorer as an administrator or non-admin user. Your installation options during automatically initiated installation procedures depend on the capabilities of the user that you are currently logged in as. IBM recommends that you install Watson Explorer and related modules as an administrator whenever possible. To do so, you should be logged in as the root user on a Linux system or as a user with Administrator privileges on a Microsoft Windows system. To eliminate the chance of installation problems related to running the Watson Explorer installation process as a user with incorrect privileges, the instructions in this section do not use an

automatically-initiated installation process, but require that you explicitly initiate installation as a user with specific privileges.

To begin installing Watson Explorer and related modules using the graphical Installation Manager interface, do the following:

1. Change directory to the location of your installation media: either the DVD device on your system or the directory into which you extracted the contents of an installation archive.
2. Launch the graphical installer by doing one of the following:
 - (Linux, Administrator installation)
 - Start a terminal session (**xterm**, **gnome-terminal**, or other terminal application) and use the **su** or **sudo -s** command to become the root user.
 - Execute the following command:

```
sh ./launchpad.sh
```
 - (Linux, Non-admin installation)
 - Start a terminal session (**xterm**, **gnome-terminal**, or other terminal application) as a normal, non-admin user.
 - Execute the following command:

```
sh ./launchpad.sh
```
 - (Windows, Administrator installation)
 - Right-click the **launchpad64.exe** application
 - Select **Run as administrator** from the context-sensitive menu that displays
 - (Windows, Non-admin installation)
 - Double-click the **launchpad64.exe** application
3. The main Watson Explorer installation screen displays. Select **Admin Install** if you are installing Watson Explorer and related modules as an administrator. Select **user Install** if you are installing Watson Explorer in non-admin mode.
4. An initial **Installation Packages** screen displays that enables you to specify the version of Watson Explorer and related modules that you are installing. Click **Next** to proceed.
5. The **Read the following license agreements** screen displays. Scroll through the complete text of the license agreement, then select the **I accept the terms in the license agreements** radio button in order to proceed with the installation process. To exit the installation process without installing Watson Explorer, select the **I do not accept the terms in the license agreement**. Click **Next** to proceed. If you accepted the license agreement, you will proceed to the next step of the installation process. If you did not accept the terms of the license agreement, Installation Manager exits.
6. A **Select a location for shared resources directory and a location for Installation Manager** screen displays that enables you to specify:
 - the directory in which shared resources that can be used by multiple packages that have been installed by Installation Manager are stored. To change this directory, modify the default directory (`/opt/IBM/IBMShared` on Linux systems or `\Program Files\IBM\IBMShared` on Microsoft Windows systems) that is displayed or browse to a new directory location.
 - the directory in which Installation Manager will be installed are stored. To change this directory, modify the default directory (`/opt/IBM/InstallationManager/eclipse` that is displayed or browse to a new directory location.

If you do not want to change these entries, or after you have made any changes that you want to make to these values, click **Next** to proceed.

7. A **Package Group** selection screen displays. A package group is a logical set of one or more compatible packages that can share resources and a common install path. The name of the Watson Explorer package group, **IBM Watson Explorer**, is preset and cannot be changed, though you can modify the name of the directory in which the package group is installed. To change this directory, modify the default directory (/opt/ibm/IDE on Linux systems or \Program Files\IBM\IDE on Microsoft Windows systems) that is displayed or browse to a new directory location. Click **Next** to proceed.
8. A **Select the features to install** screen displays. This screen lists the Watson Explorer modules that are available for installation. Click the checkbox beside each module that you want to install. Click **Next** to proceed with the installation process.
9. Depending on the features that you selected for installation, some or all of the following screens will display:

Note: No configuration panel is displayed for BigIndex, regardless of whether you selected it for installation, because it only provides an API and associated libraries that require no configuration.

configuration panel

- (Linux Only) **Installation user and group** - enables you to specify the Linux user and group that will own your Watson Explorer installation after it is installed. This user and group are typically the user and group as whom your system's web server executes. See "Configuring the Installation User and Group" on page 18 for more information about the options on this screen. Click **Next** to proceed. If you are installing Watson Explorer Engine in non-admin mode, this screen collects information about the user and group that Watson Explorer Engine should run as.
- If you selected Application Builder for installation, the **Application Builder** screen displays, which enables you to specify configuration settings related to Application Builder and application authentication. See "Configuring Application Builder" on page 18 for more information about the options on this screen. Click **Next** to proceed.
- If you selected Watson Explorer Engine for installation, the **Watson Explorer Engine Embedded Web Server** screen displays, which enables you to specify configuration settings for the embedded web server that is included with Watson Explorer Engine. See "Configuring the Watson Explorer Engine Embedded Web Server" on page 20 for more information about the options on this screen. Click **Next** to proceed.
- If you selected Knowledge Center for installation, the **Knowledge Center** screen displays, which enables you to specify configuration settings for the Knowledge Center that provides online help for Watson Explorer Engine and Application Builder. See "Configuring the Knowledge Center" on page 20 for more information about the options on this screen. Click **Next** to proceed.

If you did not select the Knowledge Center for installation, but are installing either Watson Explorer Engine or Application Builder, a screen displays that enables you to customize the URL that those applications will use to access online help.

- If you selected Results Module for installation, the **Results Module** screen displays, which enables you to specify configuration settings for the Watson

Explorer Results Module. See “Configuring the Results Module” on page 21 for more information about the options on this screen. Click **Next** to proceed.

- If you selected ZooKeeper for installation, the **ZooKeeper** screen displays, which enables you to specify configuration settings for ZooKeeper. See “Configuring ZooKeeper” on page 22 for more information about the options on this screen. Click **Next** to proceed.
10. Once you have completed any feature-specific screens, a **Summary** screen displays, showing summary information about the applications and modules that you have selected for installation, and any associated custom information that you have specified. After reviewing the entries on this screen, click **Install** to begin installing the selected features.
 11. Once the installation completes, a summary screen that lists the installed modules displays. Click **Finish** to exit Installation Manager, close its window, and return to your browser.

Once you've finished installing the Watson Explorer module(s) that you selected in the installation process, some additional steps may be required for system and inter-module integration. Proceed to “Performing Post-Installation Configuration” on page 35 for the appropriate information for the module(s) that you selected.

Configuring the Installation User and Group

Displayed when installing any Watson Explorer Engine or related modules in administrator mode on a Linux system, this panel enables you to specify the user and group that the selected modules are being installed as. The user and group ownership of all files associated with the modules that are selected for installation will be set to the specified values.

Important: The user that you specify must be a member of the group that you specify. See the documentation for your Linux distribution on the **groupmod** command (to add a user to a group, if necessary), or on the **useradd** command (which you can use to create a dedicated user in the web server group as whom to run Watson Explorer).

Configuring Application Builder

The Application Builder configuration panel enables you to specify the following during the installation process:

- **Services:**
 - **Skip Setting Service** (available only when you install in admin mode) - selecting this check box means that the service that automatically starts and stops this module will not be installed. If you do not install the service, you will always need to remember to manually start Application Builder after starting ZooKeeper, and manually stop Application Builder before stopping ZooKeeper.

Installing the service is recommended unless you have a specific reason for not doing so. If you do not install the service now, you cannot use Installation Manager to install it separately. You will have to manually create a similar service, or reinstall the module to install this service. For more information about starting and stopping services, see “Integrating Watson Explorer Into the System Startup and Shutdown Process” on page 75.
 - **Service Name** - the name of the service that will be created for use in starting and stopping Application Builder
 - **Service User** - on Linux systems, the name of the user as whom the specified service will execute

- **Display Name** - on Windows systems, the name of the service in the graphical interface.
- **Configuration:**
 - **HTTP Port** - the number of the port on which Application Builder will listen for incoming connections
 - **Server Name** - the virtual directory name that is associated with Application Builder. Incoming connections will have to be made to the URL *hostname:HTTP-Port/Server-name*.
 - **Zookeeper namespace** - the namespace that will be associated with Application Builder on any ZooKeeper installations that are configured for use by Application Builder to store configuration data
 - **Zookeeper Instance(s)** - a comma-separated list of *host:port* values that identify the hosts on which ZooKeeper has been installed for use by Application Builder and the port on which ZooKeeper is listening on each host

Note: If you are installing ZooKeeper, its configuration panel is displayed later in the installation process. If you plan to modify the ports on which your ZooKeeper instance(s) will be listening, you must identify those ports in the *host:port* entries that you enter for this field.

- **LDAP Configuration** - See your LDAP or Active Directory administrator for more information about any of the fields in this section
 - **Add LDAP Configuration** - selecting this check box indicates that you want to define a LDAP configuration for use by Application Builder, enabling you to specify values in the other fields in this part of the Application Builder configuration panel
 - **Enable SSL** - selecting this check box indicates that you want to enable the SSL security protocol
 - **LDAP Hostname** - name or IP address of the host on which the LDAP server is running
 - **LDAP Port** - port on which the LDAP service is listening on the *LDAP Hostname*
 - **LDAP Ignore Case** - check box indicating whether LDAP lookups will ignore case (checked) or be case-sensitive (un-checked)
 - **LDAP User Filter** - a filter that searches the user registry for users, enabling you to refine the LDAP users that will have access to Application Builder and BigIndex
 - **LDAP User IdMap** - specifies the LDAP filter that maps the short name of a user to an LDAP entry
 - **LDAP BaseDN** - the level of the LDAP hierarchy on the specified LDAP host where user information is located Examples: *CN=Users,DC=example,DC=com* or *OU=Users,DC=example,DC=com*
 - **LDAP BindDN** - a fully qualified account in LDAP format that has the privileges required to search for users in the LDAP hierarchy. Example: *CN=Administrator,CN=Users,DC=example,DC=com* or *OU=Administrator,DC=example,DC=com*

Note: If you enable the **Add LDAP Configuration** check box when installing Application Builder, the information that you specify on this panel will be written to the configuration file for the application server used by Application Builder

(WebSphere Liberty Profile), but will be commented out. See “Configuring Application Builder” on page 42 for information about enabling and using this LDAP configuration information.

Installing Application Builder in admin mode also installs a startup script (Linux) or service (Windows) that will automatically start Application Builder whenever you restart a system on which it is installed. This script/service is not enabled by default. See “Integrating Other Watson Explorer Modules Into the Startup/Shutdown Process” on page 78 for information about integrating Application Builder into your system's startup and shutdown processes.

Important: If you install Application Builder, make sure that the port that it uses is not blocked by a network firewall on any system where Application Builder is installed. This is the most common cause of being unable to contact Application Builder after you have started it.

Configuring the Watson Explorer Engine Embedded Web Server

The Embedded Web Server configuration panel for Watson Explorer Engine installations enables you to specify the following during the installation process:

- **Enable Embedded Webserver** - this checkbox enables, but does not start, the embedded webserver that is provided with Watson Explorer Engine
- **Port** - if you have selected the **Enable Embedded Webserver** checkbox, you can also specify the port on which the embedded webserver will listen for incoming requests

The primary cases in which you would want to enable and use the embedded webserver are if the system on which you are installing Watson Explorer Engine does not run another web server, or if you do not want to integrate Watson Explorer Engine with your system web server for some reason.

If you installed Watson Explorer Engine in admin mode, see “Integrating Watson Explorer Engine Into the Startup/Shutdown Process” on page 75 for information about integrating Watson Explorer Engine into your system's startup and shutdown processes.

Important: If you enable Watson Explorer Engine's embedded web server, make sure that the port that it uses is not blocked by a network firewall on any system where Watson Explorer Engine is installed. This is the most common cause of being unable to contact Watson Explorer Engine through the embedded web server after you have started the embedded web server.

Configuring the Knowledge Center

All online help for Watson Explorer Engine and Watson Explorer Application Builder that is not provided in tooltips or automatically generated is located in a Knowledge Center. Knowledge Centers are IBM's standard, web-based container for online help. Delivering a Knowledge Center with Watson Explorer enables installations without external Internet connectivity to continue to access online help for Watson Explorer, enables searching the online documentation for all Watson Explorer modules, and can reduce installation size by enabling multiple, local Watson Explorer installations to share access to a single Knowledge Center, rather than replicating online help in each installation.

The Knowledge Center configuration panel enables you to specify the following during the installation process:

- **Services:**

- **Skip Setting Service** (available only when you install in admin mode) - selecting this check box means that the service that automatically starts and stops this module will not be installed. If you do not install the service, you will always need to remember to manually start and stop the Knowledge Center.

Installing the service is recommended unless you have a specific reason for not doing so. If you do not install the service now, you cannot use Installation Manager to install it separately. You will have to manually create a similar service, or reinstall the module to install this service. For more information about starting and stopping services, see “Integrating Watson Explorer Into the System Startup and Shutdown Process” on page 75.

- **Service Name** - the name of the service that will be created for use in starting and stopping the Knowledge Center
- **Service User** - the name of the user as whom the specified service will execute
- **Display Name** - on Windows systems, the name of the service in the graphical interface.
- **Configuration:**
 - **Host** - the name or IP address of the host on which the Knowledge Center is running. This value is used by Watson Explorer Engine and Application Builder to locate the Knowledge Center in which their online help is located.

Note: IBM recommends that you always specify a fully-qualified domain name (or an IP address) when identifying the host on which the Knowledge Center will be running. Among other potential problems, specifying short names can cause name resolution issues if you are using a response file to install a Knowledge Center on multiple machines, which may have different default domains or use DNS differently.

- **Port** - the number of the port on which the Knowledge Center is listening for incoming connections. If you change the default port, do not use 8080 or 8081, which are the default ports for Application Builder and Results Module, respectively.
- **Virtual Directory** - the virtual directory in the Knowledge Center where online help for Watson Explorer is located.

Installing the Knowledge Center in admin mode also installs a startup script (Linux) or service (Windows) that will automatically start the Knowledge Center whenever you restart a system on which it is installed. This script/service is not enabled by default. See “Integrating Other Watson Explorer Modules Into the Startup/Shutdown Process” on page 78 for information about integrating the Knowledge Center into your system's startup and shutdown processes.

Important: If you install the Knowledge Center, make sure that the port that it uses is not blocked by a network firewall on any system where the Knowledge Center is installed. This is the most common cause of being unable to contact the Knowledge Center after you have started it.

Configuring the Results Module

The Results Module configuration panel enables you to specify the following during the installation process:

- **Services:**
 - **Skip Setting Service** (available only when you install in admin mode) - selecting this check box means that the service that automatically starts and

stops this module will not be installed. If you do not install the service, you will always need to remember to manually start and stop Results Module.

Installing the service is recommended unless you have a specific reason for not doing so. If you do not install the service now, you cannot use Installation Manager to install it separately. You will have to manually create a similar service, or reinstall the module to install this service. For more information about starting and stopping services, see “Integrating Watson Explorer Into the System Startup and Shutdown Process” on page 75.

- **Service Name** - the name of the service that will be created for use in starting and stopping the Results Module
- **Service User** - the name of the user as whom the specified service will execute
- **Display Name** - on Windows systems, the name of the service in the graphical interface.
- **Configuration:**
 - **Results Module Port** - the number of the port on which Results Module will listen for incoming connections
 - **Results Module server name** - the virtual directory name that is associated with Results Modules. Incoming connections will have to be made to the URL *hostname:Results-Module-port/Results-Module-server-name*.

Installing the Results Module in admin mode also installs a startup script (Linux) or service (Windows) that will automatically start the Results Module whenever you restart a system on which it is installed. This script/service is not enabled by default. See “Integrating Other Watson Explorer Modules Into the Startup/Shutdown Process” on page 78 for information about integrating the Results Module into your system's startup and shutdown processes.

Important: If you install the Results Module, make sure that the port that it uses is not blocked by a network firewall on any system where the Results Module is installed. This is the most common cause of being unable to contact the Results Module after you have started it.

Configuring ZooKeeper

The ZooKeeper configuration panel enables you to specify the following during the installation process:

- **Services:**
 - **Skip Setting Service** (available only when you install in admin mode) - selecting this check box means that the service that automatically starts and stops this module will not be installed. If you do not install the service, you will always need to remember to manually start and stop ZooKeeper in the correct order.

Installing the service is recommended unless you have a specific reason for not doing so. If you do not install the service now, you cannot use Installation Manager to install it separately. You will have to manually create a similar service, or reinstall the module to install this service. For more information about starting and stopping services, see “Integrating Watson Explorer Into the System Startup and Shutdown Process” on page 75.

- **Service Name** - the name of the service that will be created for use in starting and stopping ZooKeeper
- **Service User** - the name of the user as whom the specified service will execute

- **Display Name** - on Windows systems, the name of the service in the graphical interface.
- **Configuration:**
 - **MyID** - the numeric identifier for this ZooKeeper instance
 - **Client Port** - the number of the port on which Zookeeper will listen for incoming connections

Note: If you are installing Application Builder and modified the port on which Application Builder expected to find a ZooKeeper instance on the current host when you configured your Application Builder installation, you must specify the same port for the current host in this field.

- **Data Directory** - the full path-name of the directory in which ZooKeeper stores the configuration data that it is managing
- **ZooKeeper Instances** - a semicolon-delimited list of ZooKeeper instances that are part of an ensemble. Each ZooKeeper instance is of the form `server.ID=hostname:leader-port:election-port`. The *leader-port* is the port that followers use to connect to an ensemble leader, and the *election-port* is the port that all members of an ensemble use when electing the ensemble leader.
- **Additional Configuration** - a command-delimited list of configuration options that will be set in the configuration for the ZooKeeper instance that you are installing

Installing Zookeeper in admin mode also installs a startup script (Linux) or service (Windows) that will automatically start Zookeeper whenever you restart a system on which it is installed. This script/service is not enabled by default. See “Integrating Other Watson Explorer Modules Into the Startup/Shutdown Process” on page 78 for information about integrating Zookeeper into your system's startup and shutdown processes.

Important: If you install Zookeeper, make sure that the port that it uses is not blocked by a network firewall on any system where Zookeeper is installed. This is the most common cause of any Watson Explorer module being unable to contact Zookeeper after you have started Zookeeper.

Installing Watson Explorer from the Command-Line

After mounting the physical media or extracting the contents of a downloaded installation archive that contains Watson Explorer and related modules, you can install Installation Manager and any Watson Explorer module(s) in command-line mode as described in this section.

In order to install Watson Explorer from the command-line, you must first install Installation Manager from the command-line. You can then install any of the Watson Explorer and related product modules from the command-line. The next two sections explain how to do each of these as interactive installations from the command line. For information about using response files to automate command-line installation, see “Using Response Files for Silent Installation” on page 47.

Note: The different steps and associated text in the command-line installation process are referred to as *screens* because they are discrete stages in the command-line installation process, even though they do not actually clear the screen in your terminal window.

Tip: If you accidentally type an invalid response during the command-line installation process, a screen displays that identifies the invalid entry. You can return to the installation process from this screen by typing 0 (for "OK") and pressing Enter.

Important: Typing Control-c during the command-line installation process in an attempt to terminate the installation process can confuse the command-line version of Installation Manager because of the way in which responses are captured and processed. Do not use Control-c to attempt to exit from the installer. Instead, wait for the next prompt that provides the opportunity to cancel the installation.

To install Installation Manager from the command-line:

1. If installing as an administrator:
 - (Linux) - Start a terminal session (**xterm**, **gnome-terminal**, or other terminal application) and use the **su** or **sudo -s** command to become the root user
 - (Windows) - Start a command prompt or Powershell window as Administrator by right-clicking the appropriate icon or menu command and selecting **Run as administrator**

If installing in non-admin mode:

- (Linux) - Start a terminal session (**xterm**, **gnome-terminal**, or other terminal application)
 - (Windows) - Start a command prompt or Powershell window by double-clicking the appropriate icon or selecting the appropriate menu command
2. In that window, change directory to the DVD drive or the directory that contains the extracted contents of your media archive.
 3. If installing as administrator:

- (Linux) - Execute the command **./launchpad-console.sh -m**
- (Windows) - Execute the command **launchpad-console.bat -m**

If installing in non-admin mode:

- (Linux) - Execute the command **./launchpad-console.sh -n**
- (Windows) - Execute the command **launchpad-console.bat -n**

After a copyright statement and several status messages, a screen like the following displays:

```
=====> IBM Installation Manager> Install
Select packages to install:
  1. [X] IBM(R) Installation Manager 1.6.3.1
     0. Check for Other Versions, Fixes, and Extensions
     N. Next,      C. Cancel
-----> [N]
```

Note: the value inside the square brackets (in this case, the letter N is the default response that is assumed if you simply press the Enter key on your keyboard.

4. Enter N and press Enter or simply press Enter to proceed with the installation process, or press C and press Enter to cancel the installation process and return to your terminal session.

A screen like the following displays:

```

====> IBM Installation Manager> Install> Licenses

Read the following license agreements carefully.
View a license agreement by entering the number:
  1. IBM Installation Manager - License Agreement

Options:
  A. [ ] I accept the terms in the license agreement
  D. [ ] I do not accept the terms in the license agreement

  B. Back,      C. Cancel
-----> [C]

```

5. Type 1 and press Enter to read the text of the license agreement. You can page through the license agreement or interact with the license screen by using the following keys:
- **A** - display all of the license text without pagination
 - **C** - cancel the installation process and return to the terminal window
 - **D** - display the next page of the license agreement
 - **R** - return to the previous screen (the screen on which licenses are listed)
 - **U** - display the previous page of the license agreement

After reading the license, type R to return to the previous screen.

6. You must accept the terms of the license agreement in order to install Installation Manager. Type A and press Enter to accept those terms and proceed with the installation, or do any of the following to cancel your installation:

- Press C and press Enter
- Press Enter
- Press D and press Enter

Typing C and pressing Enter, or simply pressing Enter exits the installation process immediately.

Typing D and pressing Enter displays a short message which explains that you have declined the license agreement, and that installation cannot proceed. You can type R and press Enter to accept the license agreement and proceed, or type C to exit the installation process.

If you have accepted the license agreement and are continuing with the installation process, a screen like the following displays:

```

====> IBM Installation Manager> Install> Licenses> Location

Installation Manager installation location:
/opt/IBM/InstallationManager/eclipse

Options:
  L. Change Installation Manager Installation Location

  B. Back,      N. Next,  C. Cancel
-----> [N]

```

7. This screen shows the default location to which Installation manager will be installed (/opt/IBM/InstallationManager/eclipse) and enables you to change that location by typing L and entering a new installation location.

If you do not want to change the installation location or after specifying a new one, type N followed by Return, or simply press Return to proceed. A screen like the following displays:

```

=====> IBM Installation Manager> Install> Licenses> Location> Summary

Target Location:
  Package Group Name      : IBM Installation Manager
  Installation Directory   : /opt/IBM/InstallationManager/eclipse

Packages to be installed:
  IBM® Installation Manager 1.6.3.1

Options:
  G. Generate an Installation Response File

  B. Back,      I. Install,      C. Cancel
-----> [I]

```

8. (Optional) To generate a response file that contains the responses that you have made to the prompts displayed during the Installation Manager installation process, type G and press Return to be prompted for the name of the response file that you want to create. The filename that you specify must be of the form *NAME.EXTENSION*. If the filename does not contain a period, the response file creation process may exit abnormally. The location of the file can be specified as a full or relative filename, but must be a location that can be written by the user as who you are installing the software.

Tip: Creating a response file at this point may be useful if you want to use response files to install Installation Manager on multiple systems.

After entering a filename and pressing Return, a confirmation screen like the following displays once the file has been written:

```

=====> IBM Installation Manager> Install> Licenses> Location> Summary>
  Save response file> Status

OK STATUS: Response file was generated successfully: /home/joeuser/im-response.xml

  0. OK

-----> [0]

```

Type 0 and press Return, or simply press Return to proceed. The same screen shown in the previous step re-displays.

9. To begin the installation process, type I and press Return, or simply press Return. The installation process displays a histogram and a series of dots as the installation proceeds. When the installation process completes, a screen like the following displays:

```

=====> IBM Installation Manager> Install> Licenses> Location> Summary>
  Completion

The install completed successfully.

Options:
  R. Restart Installation Manager
-----> [R]

```

10. Type R and press Return, or simply press Return to restart Installation Manager and proceed. A screen like the following displays:

```

====> IBM Installation Manager

Select:
  1. Install - Install software packages
  2. Update - Find and install updates and fixes to installed software packages
  3. Modify - Change installed software packages
  4. Roll Back - Revert to an earlier version of installed software packages
  5. Uninstall - Remove installed software packages

Other Options:
  L. View Logs
  S. View Installation History
  V. View Installed Packages
-----
  P. Preferences
-----
  A. About IBM Installation Manager
-----
  X. Exit Installation Manager
----->

```

11. Type X and Press Return to exit the installer and return to the prompt in your terminal window.

After Installation Manager has been installed on your system, you can proceed with command-line installation of Watson Explorer and related Watson Explorer modules by:

1. If installing as an administrator:

- (Linux) - Start a terminal session (**xterm**, **gnome-terminal**, or other terminal application) and use the **su** or **sudo -s** command to become the root user
- (Windows) - Start a command prompt or Powershell window as Administrator by right-clicking the appropriate icon or menu command and selecting **Run as administrator**

If installing in non-admin mode:

- (Linux) - Start a terminal session (**xterm**, **gnome-terminal**, or other terminal application)
- (Windows) - Start a command prompt or Powershell window by double-clicking the appropriate icon or selecting the appropriate menu command

2. In that window, change directory to the DVD drive or the directory that contains the extracted contents of your media archive.

3. If installing as administrator:

- (Linux) - Execute the command **./launchpad-console.sh -i**
- (Windows) - Execute the command **launchpad-console.bat -i**

If installing in non-admin mode:

- (Linux) - Execute the command **./launchpad-console.sh -u**
- (Windows) - Execute the command **launchpad-console.bat -u**

After a copyright statement and several status messages, a screen like the following displays:

```

=====> IBM Installation Manager> Install

Select packages to install:
  1.  IBM Watson Explorer

      0. Check for Other Versions, Fixes, and Extensions

      N. Next,      C. Cancel
-----> [N]

```

Note: the value inside the square brackets (in this case, the letter N is the default response that is assumed if you simply press the Enter key on your keyboard.

4. Enter N and press Enter or simply press Enter to proceed with the installation process, or press C and press Enter to cancel the installation process and return to your terminal session.

A screen like the following displays:

```

=====> IBM Installation Manager> Install> Licenses

Read the following license agreements carefully.
View a license agreement by entering the number:
  1. IBM Watson Explorer - License Agreement

Options:
  A.  I accept the terms in the license agreement
  D.  I do not accept the terms in the license agreement

      B. Back,      C. Cancel
-----> [C]

```

5. Type 1 and press Enter to read the text of the license agreement. You can page through the license agreement or interact with the license screen by using the following keys:
 - **A** - display all of the license text without pagination
 - **C** - cancel the installation process and return to the terminal window
 - **D** - display the next page of the license agreement
 - **R** - return to the previous screen (the screen on which licenses are listed)
 - **U** - display the previous page of the license agreement

After reading the license, type R to return to the previous screen.

6. You must accept the terms of the license agreement in order to install Installation Manager. Type A and press Enter to accept those terms and proceed with the installation, or do any of the following to cancel your installation:
 - Press C and press Enter
 - Press Enter
 - Press D and press Enter

Typing C and pressing Enter, or simply pressing Enter exits the installation process immediately.

Typing D and pressing Enter displays a short message which explains that you have declined the license agreement, and that installation cannot proceed. You can type R and press Enter to accept the license agreement and proceed, or type C to exit the installation process.

If you have accepted the license agreement and are continuing with the installation process, a screen like the following displays:

```

====> IBM Installation Manager> Install> Licenses> Shared Directory

Shared Resources Directory:
/opt/IBM/IBMIMShared

Options:
  M. Change Shared Resources Directory

  B. Back,      N. Next,  C. Cancel
-----> [N]

```

7. This screen shows the default location to which common files that are used by multiple packages that are installed by Installation Manager will be written (/opt/IBM/IBMIMSHARED) and enables you to change that location by typing L and entering a new location for these shared resources.

If you do not want to change the installation location or after specifying a new one, type N followed by Return, or simply press Return to proceed. A screen like the following displays:

```

====> IBM Installation Manager> Install> Licenses> Shared Directory> Location

New package group:
  IBM
  Watson
  Explorer

Selected group id: "      IBM
                  Watson
                  Explorer
                  "

Selected location: "/opt/ibm/IDE"
Selected architecture: 64-bit

Options:
  M. Change Location

  B. Back,      N. Next,  C. Cancel
-----> [N]

```

8. This screen shows the default location to which the files in the Watson Explorer package group will be installed (/opt/ibm/IDE), the architecture associated with the files in that package group, and enables you to change the installation location by typing M and entering a new installation location.

If you do not want to change the installation location or after specifying a new one, type N followed by Return, or simply press Return to proceed. A screen like the following displays:

```

====> IBM Installation Manager> Install> Licenses> Shared Directory>
Location> Features

IBM Watson Explorer
  1. [ ] Application Builder
  2. [ ] BigIndex
  3. [ ] Engine
  4. [ ] Knowledge Center
  5. [ ] Results Module
  6. [ ] ZooKeeper

  B. Back,      C. Cancel
-----> [C]

```

9. Type the numbers corresponding to the modules that you want to install, pressing Return after each. You cannot enter multiple package numbers on a line. After entering all such numbers, type N and press Return, or simply press Return to proceed.
10. If you are not installing Watson Explorer on a Linux system as an Administrator (in other words, as the **root** user), skip to step “Installing Watson Explorer from the Command-Line” on page 23.

If you are installing Watson Explorer on a Linux system as an Administrator and the user that you originally logged in as is not a member of the Linux group that your web server runs as, the screen displayed by the previous step in the command-line installation process will begin with a section like the following:

```
=====> IBM Installation Manager> Install> Licenses> Shared Directory>
      Location> Features> Configurations

Configuration for IBM Watson Explorer
  1. [Error] Installation User
     Other options

      B. Back,      C. Cancel
-----> [1]
```

11. Type 1 and press Return to see the screen that enables you to change the user (or group) as whom Watson Explorer and related Watson Explorer modules are being installed. A screen like the following displays:

```
---- IBM Watson Explorer > Installation User

Watson Explorer User      : joeuser
Watson Explorer Group    : apache

Other Options:
  1. User Watson Explorer will be installed and run as
  2. Group Watson Explorer will be installed and run as

ERROR: Multiple errors occurred.

      V. View Message Details

      B. Back,      R. Return to Configurations Menu,      C. Cancel
-----> [R]
```

Note: If you are installing in non-admin mode, a similar screen displays that enables you to specify information about the user and group that Watson Explorer Engine will execute as.

12. Compare the information about the user and group that your web server runs as (which you collected in “Collecting User/Group Information for Privileged Linux Installations” on page 10, and correct the values that are displayed, as needed.

To change the user that Data Explorer will be installed and will run as, type 1, press Return, enter the new user when prompted to do so, and press Return.

To change the group that Watson Explorer will be installed and will run as, type 2, press Return, enter the new group when prompted to do so, and press Return.

type R and press Return, or simply press Return to proceed to the next configuration screen.

13. Depending on the modules that you selected for installation, some or all of the following screens will display:

- If you selected Application Builder for installation, the **Application Builder** screen displays, which enables you to specify configuration settings related to Application Builder and application authentication, and will look much like the following screen:

```
---- IBM Watson Explorer > Application Builder

Services
Service Name      : appbuilder-service-default
Service User      : joeuser

Configuration
HTTP Port         : 8080
Server Name       : AppBuilder
ZooKeeper namespace : production

ZooKeeper Instance(s) is a comma(,) delimited list
ZooKeeper Instance(s) : localhost:2181

LDAP Configuration
  1. [ ] Add LDAP configuration

Other Options:
  2. Set Application Builder service name
  3. Set Application Builder service user
  4. Set HTTP port
  5. Set server name
  6. Set ZooKeeper namespace
  7. Set ZooKeeper instance(s)

B. Back,   N. Next,   R. Return to Configurations Menu, C. Cancel
```

Note: The **Services** portion of this screen is not displayed when installing in non-admin mode.

See “Configuring Application Builder” on page 18 for more information about the options on this screen.

To modify any of the items on this screen, enter the number to the left of the item that you want to modify and press Return. You will be prompted for the new value that you want to supply. After entering a new value, press Return to accept that value and return to the configuration screen. When you are done modifying values on this screen, type N and press Return to proceed to the next configuration screen.

- If you selected Watson Explorer Engine for installation, the **Watson Explorer Engine Embedded Web Server** screen displays, which enables you to specify configuration settings for the embedded web server that is included with Watson Explorer Engine, and will look much like the following screen:

```
---- IBM Watson Explorer > Engine

  1. [ ] Enable Embedded Webserver

B. Back,   N. Next,   R. Return to Configurations Menu, C. Cancel
-----> [N]
```

See “Configuring the Watson Explorer Engine Embedded Web Server” on page 20 for more information about the options on this screen.

To modify any of the items on this screen, enter the number to the left of the item that you want to modify and press Return. You will be prompted

for the new value that you want to supply. After entering a new value, press Return to accept that value and return to the configuration screen. When you are done modifying values on this screen, type N and press Return to proceed to the next configuration screen.

- If you selected Knowledge Center for installation, the **Knowledge Center** screen displays, which enables you to specify configuration settings for the Knowledge Center that provides online help for Watson Explorer Engine and Application Builder, and will look much like the following screen:

```
---- IBM Watson Explorer > Knowledge Center

Services
Service Name      : knowledgecenter-service-default
Service User      : joeuser

Configuration
Hostname          : localhost.localdomain
Port              : 9999
Virtual Directory : help

Other Options:
  1. Set Knowledge Center service name
  2. Set Knowledge Center service user
  3. Set Knowledge Center hostname
  4. Set Knowledge Center port
  5. Set Knowledge Center virtual directory

      B. Back,   N. Next,   R. Return to Configurations Menu, C. Cancel
-----> [N]
```

Note: The **Services** portion of this screen is not displayed when installing in non-admin mode.

See “Configuring the Knowledge Center” on page 20 for more information about the options on this screen.

To modify any of the items on this screen, enter the number to the left of the item that you want to modify and press Return. You will be prompted for the new value that you want to supply. After entering a new value, press Return to accept that value and return to the configuration screen. When you are done modifying values on this screen, type N and press Return to proceed to the next configuration screen.

- If you selected Results Module for installation, the **Results Module** screen displays, which enables you to specify configuration settings for the Watson Explorer Results Module, and will look much like the following screen:

```
---- IBM Watson Explorer > Results Module

Services
Service Name      : resultsmodule-service-default
Service User      : joeuser

Configuration
Results Module port      : 8081
Results Module server name : ResultsModule

Other Options:
  1. Set Results Module service name
  2. User the Results Module service will run as
  3. Set Results Module port
  4. Set Results Module server name

      B. Back,   N. Next,   R. Return to Configurations Menu, C. Cancel
-----> [N]
```

Note: The **Services** portion of this screen is not displayed when installing in non-admin mode.

See “Configuring the Results Module” on page 21 for more information about the options on this screen.

To modify any of the items on this screen, enter the number to the left of the item that you want to modify and press Return. You will be prompted for the new value that you want to supply. After entering a new value, press Return to accept that value and return to the configuration screen. When you are done modifying values on this screen, type N and press Return to proceed to the next configuration screen.

- If you selected ZooKeeper for installation, the **ZooKeeper** screen displays, which enables you to specify configuration settings for ZooKeeper, and will look much like the following screen:

```
---- IBM Watson Explorer > ZooKeeper

Services
Service Name      : zookeeper-service-default
Service User      : joeuser

Configuration
MyID              : 1
Client Port       : 2181
Data Directory    : /opt/ibm/IDE/ZooKeeperData

Instances and Additional Configuration are semicolon(;) delimited lists
ZooKeeper Instances : server.1=localhost:16777:16888
Additional Configuration : initLimit=1000;syncLimit=1000;tickTime=2000;...

Other Options:
 1. Set service name
 2. Set ZooKeeper service user
 3. Set ID of this ZooKeeper instance
 4. Set client port
 5. Set directory for storing snapshots
 6. Add ZooKeeper instances. Entries are semicolon(;) delimited.
 7. Add additional configuration. Entries are semicolon(;) delimited.

      B. Back,   N. Next,   R. Return to Configurations Menu, C. Cancel
-----> [N]
```

Note: The **Services** portion of this screen is not displayed when installing in non-admin mode.

See “Configuring ZooKeeper” on page 22 for more information about the options on this screen.

To modify any of the items on this screen, enter the number to the left of the item that you want to modify and press Return. You will be prompted for the new value that you want to supply. After entering a new value, press Return to accept that value and return to the configuration screen. When you are done modifying values on this screen, type N and press Return to proceed to the next configuration screen.

14. Once you have completed any feature-specific screens, a **Summary** screen displays, showing summary information about your Watson Explorer installation, which will look like the following screen:

```

====> IBM Installation Manager> Install> Licenses> Shared Directory>
Location> Features> Configurations> Summary

Target Location:
Package Group Name      :          IBM
  Watson
  Explorer

Installation Directory  :    /opt/ibm/IDE
Shared Resources Directory :  /opt/IBM/IBMIMShared

Translations:
English

Packages to be installed:
IBM Watson Explorer

Options:
  G. Generate an Installation Response File

  B. Back,   I. Install,   C. Cancel

-----> [I]

```

15. If you want to create a response file containing the configuration information that you provided up to this point, type G and press Return. You will be prompted for the full pathname of the response file that you want to create. After entering that path and filename, press Return to create the specified file. After the response file is created, you will see the following screen:

```

====> IBM Installation Manager> Install> Licenses> Shared Directory>
Location> Features> Configurations> Summary> Save response file> Status

OK STATUS: Response file was generated successfully: /home/joeuser/ide-response.xml

  0. OK

-----> [0]

```

16. Type o and press **Return**, or simply press Return to re-display the summary screen again.
17. After reviewing the entries on the summary screen, type I and press **Return**, or simply press Return to begin installing the selected feature(s).
18. Once the installation completes, a summary screen that lists the installed modules displays. Type F and press Return (or simply press Return to exit Installation Manager and return to your command prompt.

Once you have finished installing the Watson Explorer module(s) that you selected in the installation process, some additional steps may be required for system and inter-module integration. Proceed to “Performing Post-Installation Configuration” on page 35 for the appropriate information for the module(s) that you selected.

If you installed Watson Explorer Engine in admin mode, see “Integrating Watson Explorer Engine Into the Startup/Shutdown Process” on page 75 for information about integrating Watson Explorer Engine into your system's startup and shutdown processes.

Installing Application Builder, the Results Module, or ZooKeeper in admin mode also installs a startup script (Linux) or service (Windows) for each module that will automatically start that module whenever you restart a system on which it is installed. This script/service is not enabled by default. See “Integrating Other

Watson Explorer Modules Into the Startup/Shutdown Process” on page 78 for information about integrating these modules into your system’s startup and shutdown processes.

Important: After installing any Watson Explorer module, make sure that the port that it uses is not blocked by a network firewall on any system where that module is installed. This is the most common cause of being unable to contact any Watson Explorer module after you have started it.

Performing Post-Installation Configuration

After installing any Watson Explorer modules, each of those modules requires some manual configuration in order to enable modules such as the Watson Explorer Engine and Results Module to interoperate and to correctly integrate the module(s) that you installed into your system environment. The next few sections explain configuration steps that are required for various Watson Explorer modules and combinations of those modules.

Important: One especially important aspect of completing any Watson Explorer installation is to integrate the services and startup scripts for Watson Explorer and related software modules into your system’s startup and shutdown processes. You can only do this if you have installed Watson Explorer and related modules as an administrator or the root user, which should be the case for all actual deployments of the Watson Explorer software. See “Integrating Watson Explorer Into the System Startup and Shutdown Process” on page 75 for more information.

Enabling Javascript in Internet Explorer 7 and Later About this task

As discussed in the **Client Requirements** section of the **Watson Explorer Engine User Manual**, a browser must support Javascript in order to use the Watson Explorer Engine administration tool. To enable Javascript in IE 7.0 and later, do the following:

Procedure

1. Select the Internet Explorer **Tools** menu, click **Internet Options**, and click the **Security** tab.
2. Select the **Internet** zone.
3. Click **Custom level...** The Security Settings - Internet Zone dialog displays.
4. In the Security Settings - Internet Zone dialog, scroll down to the **Scripting** section and click **Enable** under the **Active scripting** entry.
5. Click **OK** to accept your changes and close the Security Settings - Internet Zone dialog.
6. Click **OK** to close the Internet Options dialog.

Results

If you had tried to visit a page that requires Javascript before enabling support for active scripting, click the **Refresh** button in Internet Explorer to revisit the page and execute the Javascript that it contains.

Final Watson Explorer Engine Configuration

The next sections discuss platform-specific configuration requirements for starting and using Watson Explorer Engine and the Watson Explorer Engine administration tool on supported platforms.

Post-Install Configuration on Linux:

Note: If you have installed Watson Explorer Engine in non-admin mode, you will almost certainly want to access Watson Explorer Engine through its embedded web server, rather than integrating it with a system web server. If that is the case, skip this section and see “Completing the Installation Process (All Platforms)” on page 39 for information about configuring the embedded web server, starting the embedded web server for the first time, and logging in to the Watson Explorer Engine administration tool.

If you are running the Apache web server on a server that is running the Linux operating system and have installed Watson Explorer Engine in privileged mode and in its default location, the following command will integrate your Watson Explorer Engine installation into your Apache Web server installation:

```
cd /etc/httpd/conf.d
cp /opt/ibm/IDE/Engine/vivisimo-apache.conf .
```

If you are running the Apache web server on a Windows system, you can use the DOS or Powershell **copy** command and substitute the appropriate paths in these instructions to add the Watson Explorer Engine configuration data to your Apache installation.

On Linux systems, creating a symbolic link to the Apache configuration file is a more elegant solution, but can cause the Apache web server to fail if you subsequently move your Watson Explorer Engine installation, because the symbolic link will be unresolvable.

Note: If multiple versions of Apache are installed on your system, make sure that you integrate Watson Explorer Engine into the configuration of the Apache installation that you want to provide access to the Watson Explorer Engine administration tool.

By default, Watson Explorer Engine's Apache integration also assumes that Apache is configured to allow access to all virtual directories. If your Apache installation is configured with more restrictive permissions, you may need to un-comment the **Allow from all** directive in the *vivisimo-apache.conf* file before copying that file or after creating a symbolic link to it. You must then restart your web server before you can access the virtual directory that is used by the Watson Explorer Engine product.

For more detailed information about Apache configuration directives, see the Apache web server documentation.

Post-Install Configuration on Windows:

Note: If you have installed Watson Explorer Engine in non-admin mode, you will almost certainly want to access Watson Explorer Engine through its embedded web server, rather than integrating it with a system web server such as IIS. If that is the case, skip this section and see “Completing the Installation Process (All Platforms)” on page 39 for information about configuring the embedded web server, starting the embedded web server for the first time, and logging in to the Watson Explorer Administration tool.

After installing Watson Explorer Engine on a Microsoft Windows system on which you want to integrate Watson Explorer Engine with IIS, you will need to perform some additional configuration to ensure that Watson Explorer Engine will operate

correctly. The tasks that you will need to perform depend on the version of IIS that you are running, and are described in the next two sections.

Creating a Virtual Directory in an IIS 6 Web Server:

About this task

If the Windows system on which you have installed Watson Explorer Engine is running IIS 6, you will have to manually create a new virtual directory in IIS and associate that virtual directory with your Watson Explorer Engine installation.

Note: The *vivisimo* virtual directory is created automatically for Watson Explorer Engine installations that are integrated with the Apache or IIS 7 and higher web servers.

To manually create a virtual directory for Watson Explorer Engine installations on systems that are running IIS 6, do the following:

Procedure

1. Open the Internet Information Services (IIS) Manager, and navigate to the web site in which Watson Explorer Engine has been installed.
2. Create a new virtual directory by choosing **New** from the Action menu, and then clicking **Virtual Directory**.
3. A wizard will then guide you through creating a Virtual Directory. Enter the following information when prompted:
 - **Alias:** enter the new virtual directory name. The recommended alias for a manually-created virtual directory is *vivisimo*, which is the name of the virtual directory that is created automatically on non-IIS 6 web servers.
 - **Path:** enter the full path name of the *www* folder of the Watson Explorer Engine installation directory (typically *C:\Program Files\IBM\IDE\Engine\www*).
 - **Permissions:** Select **Read**, **Run Scripts**, and **Execute**.
4. Click **Finish** to save the definition for the new virtual directory.

After manually creating a virtual directory, you will need to integrate it into your Watson Explorer Engine installation if the name of the virtual directory that you created (its *Alias*) is not *vivisimo*, which is the default value for a Watson Explorer Engine virtual directory. To tell your Watson Explorer Engine installation to use a custom virtual directory:

5. Edit the file *vivisimo.conf*, which is located in the directory in which you installed Watson Explorer Engine.
6. Add a statement like the following inside the **<options>** element:

```
<option name="base-url">VIRTUAL-DIRECTORY</option>
```

Replace *VIRTUAL-DIRECTORY* with the name of the virtual directory that you created for Watson Explorer Engine, preceded and followed by a forward slash (/).

For example, if you set the name of the virtual directory for your Watson Explorer Engine installation to *search*, you would add the following entry to your *vivisimo.conf* file:

```
<option name="base-url">/search/</option>
```

7. Save the updated file and restart your Watson Explorer Engine installation to ensure that it uses the new *base-url* setting when constructing internal URLs. See “Command-Line Watson Explorer Engine Administration Commands” on page 68

page 68 for information on using the **velocity-shutdown** and **velocity-startup** commands to restart your Watson Explorer Engine installation.

Enabling CGI Module Support and Execution:

Using Watson Explorer Engine with all supported versions of the Microsoft IIS web server requires that support for unspecified CGI modules be enabled in the IIS Web Server.

Procedure

1. Start the Internet Information Services (IIS) Manager.
2. In the **Connections** pane, select your server, and then from the center pane, select **ISAPI and CGI Restrictions**.
3. From the Actions pane, click **Open Feature**.
4. From the Actions pane, click **Edit Feature Settings**.
5. Ensure that **Allow unspecified CGI modules** is selected. If it is not, select that option and click **OK** to save your changes and close the window.

What to do next

Installing and using Watson Explorer Engine with IIS also requires that execute rights be enabled for CGI applications so that your IIS web server can run executable files that are in the `www\cgi-bin` subdirectory of the directory in which you installed Watson Explorer Engine.

To enable execute rights:

- IIS 6.0 on Windows 2003 Server Systems:
 1. Select **Control Panel > Administrative Tools > Internet Information Services (IIS) Manager**
 2. Drill down to and right-click on the `vivisimo` virtual directory (or whatever virtual directory you created from “Creating a Virtual Directory in an IIS 6 Web Server” on page 37) and select **Properties**.
 3. In the **Execute Permissions** section, select **Scripts and Executables**, and click **Apply**.
 4. Depending on other web-based applications that are installed on your system, you might also need to enable `.EXE` CGI extensions. To enable these extensions:
 - a. Select **Web Service Extensions**.
 - b. If the **All unknown CGI extensions** web service extension is not marked as **Allowed**, select it and click **Allow**. A confirmation window displays, identifying this option as a potential security risk. Click **Yes** to confirm that you understand the risks and you want to enable these extensions.
- IIS 7.0 and higher on Windows 2008 and later systems:
 1. Open IIS Manager:
 - For Windows 2008 systems: Select **Control Panel > Administrative Tools > Internet Information Services (IIS) Manager**.
 - For Windows 2012 systems: From the desktop, hover your cursor in the upper- or lower- right corner of the screen to show the charms. Click the **Start** charm, and then click the **Internet Information Services (IIS) Manager** tile.

2. In the Connections pane, expand the entry for your computer system, expand **Sites**, and select the default site, which is the site in which the Watson Explorer Engine virtual directory (vivisimo) is created.
3. Double-click the **Handler Mappings** feature.
 - If the state of **CGI-exe** is **Enabled**, no additional configuration is necessary. Proceed to “Completing the Installation Process (All Platforms).”
 - If the state of **CGI-exe** is **Disabled**, select it and click **Edit Feature Permissions** from the Actions pane. The Edit Feature Permissions window displays.
Select **Execute** and click **OK** to enable applications with the .exe extension to execute in response to CGI requests. If **Execute** is disabled, select **Script > Execute**, and then click **OK**.

Important: When you install Watson Explorer Engine and integrate it with an existing Microsoft Windows IIS server, the installation process grants write permission to all users on all Watson Explorer directories that can be written to. Administrators can then further restrict write access to the Watson Explorer Engine files and directories manually.

Completing the Installation Process (All Platforms): Open the Watson Explorer Engine administration tool in a web browser. If you used the default installation locations provided by Installation Manager, installed Watson Explorer Engine in privileged mode, and integrated Watson Explorer Engine with a system web server, the Watson Explorer Engine Administration tool should be available at the following URL, where *YOURSERVER* is the name or IP address of the system on which your web server is running:

`http://YOURSERVER/vivisimo/cgi-bin/admin`

Note: You can also complete your Watson Explorer Engine installation using its embedded web server. See “Configuring Watson Explorer Engine’s Embedded Web Server” on page 61 for information about configuring the embedded webserver so that you can use it to contact the Watson Explorer Engine administration tool.

When using the embedded web server to complete your installation, you must start the embedded web server manually after it is configured. You must do this using the **velocity-webserver start** command - you cannot use the standard **velocity-startup** command until Watson Explorer Engine has been started at least once. (See “Pre-Install Configuration and Data Collection” on page 6 for a discussion of tasks that you should perform before installing Watson Explorer Engine.)

Once the Watson Explorer Engine login prompt displays in your browser, use the default login (**data-explorer-admin**) and password (**TH1nk1710**, which is **T-H-one-n-k-one-seven-one-zero**) to access the Watson Explorer Engine administration tool. Watson Explorer Engine is ready to use!

Important: After accessing the Watson Explorer Engine administration tool for the first time using the default login and password, you should **immediately** change the password for this login. To change the password for a Watson Explorer Engine account, use the **Management > Users** menu to access and modify the current settings for that user account.

Tip: At this point, you may want to create additional Watson Explorer Engine users for development purposes using the **Management > Users** menu command.

This enables other developers to access the Watson Explorer Engine administration tool while keeping your site's administrative login and password private.

If you plan to use the embedded web server that is provided with Watson Explorer Engine, you must configure it before you can start it. You can then start the web server as part of the general Watson Explorer Engine startup process by using the **velocity-startup** command), which you can then integrate into your system's startup process:

- See “Configuring Watson Explorer Engine's Embedded Web Server” on page 61 for information about configuring the embedded web server that is provided with Watson Explorer Engine.
- See “Command-Line Watson Explorer Engine Administration Commands” on page 68 for general information about integrating Watson Explorer Engine into your system startup and shutdown processes, and for detailed information about the Watson Explorer Engine commands that are provided for use in these cases.

Tip: Before creating any Watson Explorer Engine search applications or search collections, you may want to take advantage of configuration options that enable you to store Watson Explorer Engine search collection data on filesystems other than the one in which the Watson Explorer Engine software is installed. For example, if the system on which Watson Explorer Engine is installed also has access to large, high-speed RAID storage devices, you may want to install Watson Explorer Engine in its default location but use Watson Explorer Engine's **search-collections-dir** configuration variable to specify that search collections are created and stored on the high-speed, higher-capacity storage devices. See “Configuring Global Options for Watson Explorer Engine” on page 58 for more information about how and where to set the **search-collections-dir** configuration variable.

After completing the Watson Explorer Engine installation process, see the “System Administration” on page 56 section of this document for tips on integrating Watson Explorer Engine into your IT policies and procedures. If you are running automatic anti-virus or other system scanning software, see the “Integrating Watson Explorer with Anti-Virus Software” on page 80 subsection for information on excluding specific Watson Explorer Engine directories from the scanning process.

Final Watson Explorer Results Module Configuration

The IBM Watson Explorer Results Module is an application that is powered by a IBM WebSphere Liberty Profile (WLP) server. Before the Results Module can be used, the WLP server must be started.

Configuring the Results Module: You can adjust the configuration of your Results Module WLP server by editing the `server.xml` file in the `ResultsModule/wlp/usr/servers/ResultsModule` directory of your installation. The configuration information that is displayed depends on any changes that you made on the Results Module configuration screen during the installation process, but will look something like:

```
<server description="ResultModule">

  <!-- Enable features -->
  <featureManager>
    <feature>jsp-2.2</feature>
    <feature>servlet-3.0</feature>
  </featureManager>

  <httpEndpoint id="defaultHttpEndpoint"
```

```

host="*"
httpPort="8081"
httpsPort="9088" />

<applicationMonitor updateTrigger="disabled"/>
</server>

```

You can adjust this configuration in many different ways to meet your site requirements. For example, you can change the port used to access the Results Module by changing the `httpPort="8081"` line to specify another port. You can configure other aspects of the WLP server, including https-only configuration, by editing other values in this file.

You can find detailed information about configuring the WLP server in the online WebSphere 8.5 documentation.

Starting the Results Module: Installing the Results Module as an administrator creates a service to automatically start the Results Module on Windows systems, and a startup script to automatically start the Results Module on Linux systems.

To start the Results Module after installing it as an administrator:

- (Linux) - Execute the `/etc/init.d/resultsmodule-service-default` script (or whatever other name you may have assigned to the service) as the root user or via the **sudo** command to start the Results Module for the first time. You will want to integrate this script into the startup scripts for your system's default runlevel to ensure that the Results Module starts automatically each time that you restart your system. See "Startup/Shutdown Integration on Linux Systems" on page 78 for more information.
- (Windows) - Execute the `resultsmodule-service-default` service (or whatever other name you may have assigned to the service) from the Services (Local) control panel to start the Results Module for the first time. See "Startup/Shutdown Integration on Windows Systems" on page 79 for more information.

The commands to start and stop Results Module that are explained in the rest of this section are discussed in relation to the directory where these modules are installed. If you installed the Results Module as a non-admin user, the default location where it is installed is in the `ibm/IDE` subdirectory of your home directory on Linux systems, and the `IBM/IDE` subdirectory of your home directory on Microsoft Windows systems.

To start the Results Module after installing it in non-admin (un-privileged) mode, regardless of the operating system platform on which you installed it:

- Execute the `server start ResultsModule --clean` command from the `ResultsModule/wlp/bin` directory of your installation.

Accessing the Results Module: After starting the Results Module, you can access the Results Module interface at a URL like the following:

```
<hostname>:<port>/ResultsModule
```

Where `<hostname>` is the host name or IP of the Results Module server, and `<port>` is the port used for this module (8081 by default). To access the Results Module using the Bootstrap user, use the username **bootstrap** and the password **changeme**.

The Results Module is now installed and operational. For details about configuring the Results Module, see [Creating and Configuring Watson Explorer Results Module](#).

Note: There may be a significant wait period (greater than 1 minute) when accessing the Results Module for the first time after the WLP has been started. This is part of the WLP initialization process and only occurs the first time it is accessed after startup.

Useful WLP Server Commands for the Results Module:

The following additional WLP server commands may be useful:

- To stop the Results Module use `ResultsModule/wlp/bin/server stop ResultsModule` command from your installation directory.
- To check the status of the Results Module, use the `ResultsModule/wlp/bin/server status ResultsModule` command from your installation directory.

Final Watson Explorer Application Builder Configuration

The IBM Watson Explorer Application Builder is an application that is powered by a IBM WebSphere Liberty Profile (WLP) server and a ZooKeeper server. Before the Application Builder can be used, the ZooKeeper and WLP servers must be started, and in that order.

Configuring Application Builder:

Before starting Application Builder, you may want to modify its default configuration data to adjust options such as the port on which Application Builder listens, or the virtual directory if this was not set during the installation process. You can adjust the configuration of your Application Builder WLP server from the values that you set on the configuration panel during installation by editing the `server.xml` file in the `AppBuilder/wlp/usr/servers/AppBuilder` directory of your Watson Explorer installation. The following default configuration is provided when you install the Application Builder:

Important: The security model used by Application Builder is also configured in this file. However, the security model can only be configured after the user entity and administrative users have been configured in the Application Builder interface. The procedure to configure users is described in the Entities section of the Application Builder documentation. The procedure described in that section must be performed using the default login credentials, regardless of whether you specified LDAP configuration during the installation process. Any LDAP configuration entered during the installation will be present but commented out in the `server.xml` file.

```
<server description="AppBuilder">
  <webContainer useoriginalrequeststate="true"/>
  <!-- Enable features -->
  <featureManager>
    <feature>servlet-3.0</feature>
    <feature>appSecurity-1.0</feature>
  </featureManager>
  <httpEndpoint id="defaultHttpEndpoint"
    host="*"
    httpPort="8080"
    httpsPort="9088" />
  <!-- Hard-coded username and password to allow initial configuration -->
  <basicRegistry id="basic" realm="customRealm">
```

```

    <user name="data-explorer-admin" password="TH1nk1710" />
</basicRegistry>

<!-- Sample LDAP configuration -->
<!--
<ldapRegistry
  id="ldap"
  realm="SampleLdapADRealm"
  host="ldap.sample.com"
  port="389"
  ignoreCase="true"
  baseDN="CN=Users,DC=sample,DC=ibm,DC=com"
  bindDN="sample@sample.ibm.com"
  bindPassword="PASSWORD"
  userFilter="(sampleUserName=%v)"
  userIdMap="user:SampleUserName"
  ldapType="Microsoft Active Directory"
  sslEnabled="true"
  sslRef="LDAPSSLSettings">
</ldapRegistry>

<sslDefault sslRef="LDAPSSLSettings" />
<ssl id="LDAPSSLSettings" keyStoreRef="LDAPKeyStore" trustStoreRef="LDAPTrustStore" />
<keyStore id="LDAPKeyStore"
location="C:\Program Files\IBM\IDE\AppBuilder\wlp\usr\servers\AppBuilder\sslCert.jks"
type="JKS" password="PASSWORD" />
<keyStore id="LDAPTrustStore"
location="C:\Program Files\IBM\IDE\AppBuilder\wlp\usr\servers\AppBuilder\sslCert.jks"
type="JKS" password="PASSWORD" />
-->

<application type="war" id="AppBuilder"
  name="AppBuilder" location="${server.config.dir}/apps/AppBuilder">
  <application-bnd>
    <security-role name="AllAuthenticated">
      <special-subject type="ALL_AUTHENTICATED_USERS" />
    </security-role>
  </application-bnd>
</application>

<applicationMonitor updateTrigger="disabled"
  dropinsEnabled="false"/>
</server>

```

You can adjust this configuration in many different ways to meet your site requirements. For example, you can change the port used to access the Application Builder by changing the `httpPort="8080"` line to specify another port. You can configure other aspects of the WLP server, including https-only configuration, by editing other values in this file.

Note: In Watson Explorer 9.0, values for the `userFilter` and `userIdMap` fields are not requested when configuring Application Builder prior to installation. If you plan to use LDAP authentication with Application Builder, you must supply appropriate values for these fields before uncommenting and using the LDAP configuration information in this file.

You can find detailed information about configuring the WLP server in the online WebSphere 8.5 documentation.

Starting ZooKeeper and Application Builder: Services are provided to start ZooKeeper and the Application Builder on Windows systems, and startup scripts

are provided to start ZooKeeper and Application Builder on Linux systems. These scripts and services are only installed if you install the Application Builder in admin mode.

Important: Application Builder stores its configuration information in ZooKeeper. It is extremely important that you start ZooKeeper before starting Application Builder.

To start ZooKeeper and Application Builder after installing Application Builder in admin mode:

- (Linux) - On the server(s) where ZooKeeper is installed, execute the `/etc/init.d/zookeeper-service-default`. Next, on the server(s) where Application Builder is installed, execute `/etc/init.d/appbuilder-service-default` scripts (or whatever other names you may have assigned to these scripts) - in that order - as the root user or via the **sudo** command to start ZooKeeper and Application Builder for the first time. You will want to integrate these scripts into the startup scripts for your system's default runlevel, in the same sequence, to ensure that ZooKeeper and Application Builder start each time that you restart your system. You will want to integrate these scripts into the shutdown scripts for your system's default runlevel, in reverse order, to ensure that ZooKeeper and Application Builder are shut down cleanly each time that you restart your system. See "Startup/Shutdown Integration on Linux Systems" on page 78 for more information.

Important: When starting your system, ZooKeeper should be started before Application Builder, and must be started before the application or the Application Builder admin tool is accessed. When shutting down your system, Application Builder must be shut down before shutting down ZooKeeper.

- (Windows) - On the server(s) where ZooKeeper is installed, execute the `zookeeper-service-default` service. Then, on the server(s) where Application Builder is installed, execute the `appbuilder-service-default` service (or whatever other names you may have assigned to these services) from the Services (Local) control panel application to start ZooKeeper and the Application Builder for the first time. You should start these services in that order when starting them for the first time or after restarting your system, and ZooKeeper must be started before the application or the Application Builder admin tool is accessed. You must stop these services in the reverse order (`appbuilder-service-default` and then `zookeeper-service-default`) when shutting down Application Builder. When installed on the same system, these scripts are installed so that the associated services will stop and start in the right order, each time that you shut down and restart your system. See "Startup/Shutdown Integration on Windows Systems" on page 79 for more information.

The commands to start and stop ZooKeeper and Application Builder that are explained in the rest of this section are discussed in relation to the directory where these modules are installed. If you installed Application Builder and ZooKeeper as a non-admin user, the default location where those modules are installed is in the `ibm/IDE` subdirectory of your home directory on Linux systems, and the `IBM/IDE` subdirectory of your home directory on Microsoft Windows systems.

To start ZooKeeper and Application Builder after installing Application Builder in non-admin mode:

- Start the ZooKeeper server before starting Application Builder. To start the ZooKeeper server on a Linux system, use the `Zookeeper/zookeeper/bin/`

zkServer.sh restart command from your Watson Explorer installation directory. On Microsoft Windows systems, use the Zookeeper\zookeeper\bin\zkServer.cmd command.

Note: On Microsoft Windows systems, the command to start ZooKeeper is a non-terminating script. You must create a wrapper for this script if you want to deploy it as a service.

In a default installation, ZooKeeper has been configured by default to be used on a single server with the Application Builder.

- After starting the ZooKeeper server, you can start the Application Builder WLP server with its current settings. To do so, use the **./server start AppBuilder --clean** command, which is located in the AppBuilder/wlp/bin directory of your Application Builder installation. This starts the Application Builder with the configuration that is defined in the server.xml file.

Accessing Application Builder: After starting ZooKeeper and the Application Builder, you can access the Application Builder administration tool at the following URL:

```
<hostname>:<port>/AppBuilder/admin
```

Where <hostname> is the host name or IP or the Application Builder server, and <port> is the port used for this module (8080 by default). To access the Application Builder using the Bootstrap user, log in using the username **data-explorer-admin** and the password **TH1nk1710**.

The Application Builder is now installed and operational. Consult the Developing Applications section of the Application Builder documentation for additional configuration details.

Note: There may be a significant wait period (greater than 1 minute) when accessing the Application Builder or its applications for the first time after the WLP has been started. This is part of the WLP initialization process and only occurs the first time it is accessed after startup.

Useful WLP Server Commands for Application Builder:

The following additional WLP server commands may be useful:

- To stop the Application Builder WLP Server, use AppBuilder/wlp/bin/server stop AppBuilder command from your installation directory.

Note: On Microsoft Windows systems, the Application Builder service must be stopped before the associated ZooKeeper service is stopped. If the associated ZooKeeper service is stopped first, the Application Builder service will hang when an attempt to stop the service is made.

- To check the status of the Application Builder WLP Server, use AppBuilder/wlp/bin/server status AppBuilder command from your installation directory.

Troubleshooting an Installation

The Watson Explorer software installation mechanism is designed to make it easy to install any Watson Explorer module(s) successfully on your systems. However, depending upon how your web server, firewalls, and other network security mechanisms are configured, you may encounter problems when first getting started using the module(s) that you have installed. The following are some

common errors that you may see, but which can easily be resolved by adjusting the network configuration or file protections on the system where you have installed the Watson Explorer module(s):

- **Corrupt data error when trying to login to the Watson Explorer Engine administration tool immediately after installation.**
Either TCP port 80 (HTTP), TCP port 443 (HTTPS), or both are blocked. Unblock TCP port 80 or 443.
- **Could not create lock file.//data/admin.lock.**
The permissions on the `vivisimo.conf` file are set such that the web server does not have access. Change the permissions of the `vivisimo.conf` file (located in the Watson Explorer Engine installation directory).
- **HTTP 404 Error after installation**
One or more of the following issues may be the cause:
 - The **Alias** (Apache) or **Virtual Directory** (IIS) is not configured properly, or doesn't point to the correct location
 - The URL in the browser is incorrect
 - On Linux systems: AppArmor or SELinux is set to restrictive mode and is blocking the execution of the Velocity binaries. Set these to permissive or disabled mode.
- **HTTP 500 Error when installing in Windows under IIS**
CGI executables haven't been configured to run correctly, follow the instructions for your version of Windows to allow CGI executables to run.
- **HTTP 403 Error**
The user that is logged in, or is running the web service, has insufficient privileges. Check the permissions on the account.

Getting Help

If you have questions or encounter problems with this installation procedure or any module of the Watson Explorer software, please file a service request using the IBM Support Portal. (Service requests are also sometimes referred to as PMRs - Problem Management Records.)

You must sign in to the IBM Support Portal to submit or access service requests. If you do not already have a login for the Support Portal, click **Register Now** and specify your email address as your login during the registration process. After registering, you must contact the Site Technical Contact (STC) for your company's account and request that your account be granted the ability to create service requests. Once your account has been authorized, you can click **Support and Download** from the top navigator menu, and select **Open a Technical Service Request** to enter your request.

Note: For more detailed information about submitting a service request, see the IBM Service Request Quick Start page.

To take advantage of any support that is available for your national language and locale, select your region and country from the Directory of worldwide contacts page.

Using Response Files for Silent Installation

As discussed in “Overview of Watson Explorer Installation Models” on page 6, installing any Watson Explorer or related software module creates a response file that contains the responses to all prompts and configuration settings that were made when that module was originally installed. As described in the next few sections, these response files can be used to perform a silent installation of that same module on another system, by updating the response file appropriately and then making both the response file and your installation archive or media available on that other system.

Locating Response Files

After using Installation Manager via the `launchpad.sh` or `launchpad-console.sh` script to install Watson Explorer or any related software module on a system, the response file that captures your responses to the prompts associated with that module is saved in the directory into which that module was installed. By default, Watson Explorer modules are installed into subdirectories of the directory into which you installed Watson Explorer (by default: `/opt/ibm/IDE` on Linux systems, and `Program Files\IBM\IDE` on Microsoft Windows systems) that have the name of the associated module. These directories, the Watson Explorer or related module with which they are associated, and the associated response files are:

- **AppBuilder** - Application Builder. The response file for this module is `appBuilderResponse.xml_SOFTWARE_VERSION`.
- **BigIndex** - BigIndex API and its developer documentation. The response file for this module is `BigIndexResponse.xml_SOFTWARE_VERSION`.
- **Engine** - Watson Explorer Engine. The response file for this module is `engineResponse.xml_SOFTWARE_VERSION`.
- **KnowledgeCenter** - IBM Knowledge Center, which contains online documentation for Application Builder, Watson Explorer Engine, and general-purpose documentation such as this installation guide. The response file for this module is `knowledgeCenterResponse.xml_SOFTWARE_VERSION`.
- **ResultsModule** - Results Module. The response file for this module is `resultsModuleResponse.xml_SOFTWARE_VERSION`.
- **ZooKeeper** - ZooKeeper. The response file for this module is `zooKeeperResponse.xml_SOFTWARE_VERSION`.

In each of the response files that are generated during a graphical installation of Watson Explorer, the `SOFTWARE_VERSION` portion of the file name is the Watson Explorer version with which they are associated, such as `9.0.0`.

Modifying Response Files

To use an automatically generated response file to install a Watson Explorer or related software module on another system, you must modify that response file to identify the location at which your installation media will be available on the remote machine. The value in the response file that identifies the directory in your installation media which contains the repository from which software is installed is the **location** attribute of the `<repository>` element, which is located in the `<server>` element. Each automatically generated response file that is produced during a graphical installation of Watson Explorer contains a template element of the following form:

```
<server>
  <repository location='REPOSITORY_LOCATION' />
</server>
```

Before using a response file, you must replace *REPOSITORY_LOCATION* with the full pathname to the ide-offering directory on the installation media.

Installing Using a Response File

Once you have updated the response file appropriately, you must make it and your installation media available on the remote system, which is often done by mounting a directory containing them on the remote machine using a remote filesystem protocol such as NFS (Linux) or CIFS/SMB (Windows). Once the response file and your installation media are available on the remote machine, you can use that response file for installing the associated Watson Explorer module.

As an example, suppose that you are using a NFS to make a directory containing your installation media for Watson Explorer Enterprise Edition, version 9.0.0 and the response file for the Application Builder available as the directory `/mnt/NFS/DataExplorer` on each system where you want to install Watson Explorer. In that case, the value of the `<repository>` element's *REPOSITORY_LOCATION* attribute in your response file would be `/mnt/NFS/DataExplorer/IDE_V9.0.0-EE/ide-offering`, as in the following example:

```
<server>
  <repository location='/mnt/NFS/DataExplorer/IDE_V9.0.0-EE/ide-offering' />
</server>
```

You could then do a silent installation of Application Builder on the remote system by executing the following command as the root user (or via the **sudo** command) on each system where that directory is mounted via NFS:

```
/mnt/NFS/DataExplorer/IDE_V9.0.0-EE/IM/IMLinux/tools/imcl -s
  -acceptLicense -input
  /mnt/NFS/DataExplorer/appBuilderResponse.xml_9.0.0
```

This command should be typed on a single line, and is only shown on multiple lines here for formatting purposes.

Modifying Watson Explorer Installations

Modifying an existing Watson Explorer installation means adding or removing modules that are part of the same release of Watson Explorer.

The Watson Explorer 9.0.0.1, 9.0.0.2, and 9.0.0.3 fix packs do not support modifying an existing Watson Explorer installation. To add Watson Explorer modules to an existing installation, you must shut down all Watson Explorer services, uninstall all Watson Explorer modules, and reinstall Watson Explorer, selecting the modules that you want to install.

Important: When uninstalling Watson Explorer modules, any Watson Explorer Engine search projects, sources, search collections, or Results Module and Application Builder customizations that you created are preserved. They will therefore be available in any of the modules that you reinstall.

Note: Watson Explorer installations cannot be updated (which means installing an incrementally updated version of the software) or rolled back (which means to revert to a previous version of the software) using Installation Manager. Selecting the Installation Manager controls that are associated with those operations will have no effect.

Reinstalling Watson Explorer

If you have previously installed an evaluation or beta version of Watson Explorer, you will eventually want to install an official release of Watson Explorer. Uninstalling Watson Explorer is discussed in “Uninstalling Watson Explorer Modules” on page 50. However, before re-installing Watson Explorer or related software modules using a locally-installed version of Installation Manager, you should make sure that no cached files from previous Watson Explorer installations are still located on your system. If cached files from previous Watson Explorer installations are still present on your system, version number collisions may cause out-dated files to be installed on your system.

To disable using (and retaining) cached files from previous versions of Watson Explorer:

- When using Installation Manager in graphical mode:
 1. Start the graphical version of Installation Manager and select **File > Preferences**
 2. Select **Files for Rollback** in the left pane
 3. Deselect the **Save files for rollback** checkbox
 4. Click **Delete Saved Files**
 5. Click **Apply**
 6. Click **OK** to close the Preferences dialog
- When using Installation Manager in command-line mode:
 1. Start the command-line version of Installation Manager, as discussed in “Installing Watson Explorer from the Command-Line” on page 23.
 2. Type P and press Enter to display the preferences screen
 3. Type 3 and press Enter to select the **Files for rollback** option
 4. Enter 1 to de-select the **Save files for rollback** option and press Enter
 5. If the **Delete saved files option** is listed, enter D to select that option and press Enter. (The **Save files for rollback** option is not listed if you do not currently have any saved files.)
 6. Type A to select the **Apply Changes and Return to Preferences Menu** option, and press Enter
 7. Type R and press Enter to return to the main Installation Manager screen.

Once you have configured Installation Manager so that it does not preserve these files, click **Install** on the main Installation Manager screen to begin installing Watson Explorer and related software modules as discussed in:

- Graphical installation - See “Installing Watson Explorer Graphically” on page 15. The only difference between initialization as described in that topic and reinstallation is that you will not need to execute the **launchpad.sh** (Linux) or **launchpad.bat** (Windows) scripts, because Installation Manager is already installed.
- Command-Line installation - See “Installing Watson Explorer from the Command-Line” on page 23. The only difference between initialization as described in that topic and reinstallation is that you will not need to execute the **launchpad-console.sh** (Linux) or **launchpad-console.bat** (Windows) scripts, because Installation Manager is already installed.

Uninstalling Watson Explorer and Installation Manager

In addition to installing and modifying Watson Explorer modules, Installation Manager can also be used to uninstall existing Watson Explorer installations.

Important: You can only uninstall Watson Explorer modules by executing a version of Installation Manager that is actually installed on your system. You cannot uninstall a Watson Explorer installation by using the version of Installation Manager that is provided on an ISO image, DVD, or complete downloaded Watson Explorer product archive.

Uninstalling Watson Explorer Modules

Important: Before uninstalling any Watson Explorer installation, any services that are directly associated with the installed modules must be stopped. This does not include system web servers with which Watson Explorer Engine has been integrated. If Application Builder is installed, you must stop Application Builder before stopping or otherwise terminating ZooKeeper.

Note: Watson Explorer installations cannot be updated (which means installing an incrementally updated version of the software) or rolled back (which means to revert to a previous version of the software) using Installation Manager. Selecting the Installation Manager controls that are associated with those operations will have no effect.

To uninstall one or more Watson Explorer modules, execute a version of Installation Manager that is locally installed on your system, and click **Uninstall** on its initial dialog. Subsequent dialogs will enable you to select the module(s) that you want to uninstall and then initiate the uninstallation process.

Warning: Do not attempt to uninstall any Watson Explorer modules by simply removing the directory that contains them. Installation Manager maintains its own records of packages that it has been used to install, and you may not be able to reinstall packages that were originally installed by Installation Manager but which you have manually removed.

Tip: As a convenience, Watson Explorer creates a response file called `uninstallResponse.xml` in your Watson Explorer installation directory when Watson Explorer is installed. You can use this response file to do a silent uninstall of all Watson Explorer and related software modules. Using this response file will not remove Installation Manager.

Uninstalling Installation Manager

To remove Installation Manager:

- On Microsoft Windows systems, use the Control Panel's **Add/Remove Programs** or **Uninstall a program** dialogs (depending on the version of Windows that you are using). If you installed Installation Manager as an administrator, you will need to right-click on this application and select **Run as administrator** to obtain the necessary permissions to uninstall Installation Manager.
- On Linux systems:
 - On graphical systems, select the **Applications > IBM Installation manager** menu command.
 - On command-line systems:

- Execute a command like the following to identify the packages on your system that have been installed using Installation Manager:

```
/PATH/TO/InstallationManager/eclipse/tools/imcl
listInstalledPackages
```

If you installed Installation Manager as the root user (or via the **sudo** command), replace */PATH/TO* with */opt/IBM*. If you installed Installation Manager in non-admin, replace */PATH/TO* with the name of the home directory of the user as whom you installed Installation manager.

- If this command only returns a single value of the form **com.ibm.cic.agent...**, that is Installation Manager, which you can remove with a command like the following: */PATH/TO/InstallationManager/eclipse/tools/imcl uninstall com.ibm.cic.agent...*

In this command, you should replace *com.ibm.cic.agent...* with the value that was displayed by the previous **imcl** command.

Uninstalling Watson Explorer Modules using a response file

You can uninstall IBM Watson Explorer modules by using a response file.

About this task

When you install Watson Explorer, graphically or with the command-line version of Installation Manager, the program creates a response file that you can use to perform silent uninstalls on the computer where the `uninstallResponse.xml` file was generated, and on other computers that have an identical installation..

If you installed all the modules, the `uninstallResponse.xml_versionNumber` file contains the following code:

```
<?xml version="1.0" encoding="UTF-8"?>
<agent-input>
  <uninstall>
    <offering id='com.ibm.ide.ee' profile='IBM Watson Explorer'
features='com.ibm.ide.engine.feature,com.ibm.ide.resultsmodule.feature,
com.ibm.ide.knowledgecenter.feature,com.ibm.ide.bigindex.feature,
com.ibm.ide.zookeeper.feature,com.ibm.ide.appbuilder.feature,'/>
  </uninstall>
</agent-input>
```

Procedure

From a command-line interface, enter the following command:

- On Linux systems, if you installed in admin mode:

```
opt/IBM/InstallationManager/eclipse/tools/imcl input
/opt/ibm/IDE/uninstallResponse.xml_versionNumber
```

On Linux systems, if you installed in non-admin mode:

```
userHomeDirectory/InstallationManager/eclipse/tools/imcl input
installUserHomeDirectory/IDE/uninstallResponse.xml_versionNumber
```

- On Windows systems, if you installed in admin mode, from the `C:\` directory, run:

```
"Program Files\IBM\Installation Manager\eclipse\tools\imcl" input
"C:\Program Files\IBM\IDE\uninstallResponse.xml_versionNumber"
```

On Windows systems, if you installed in non-admin mode, from the `C:\` directory, run:

```
"Users\userDirectory\IBM\Installation Manager\eclipse\tools\imcl" input
C:\Users\userDirectory\IBM\IDE\uninstallResponse.xml_versionNumber
```

Upgrading from Previous Releases

IBM makes every possible effort to ensure that software upgrades are as easy as possible. However, seamless upgrades are not always possible when product names, installation requirements, and installation locations change as they have in the transition from Velocity and IO Pro to Watson Explorer Engine and Watson Explorer Engine Results Module.

This section explains how to migrate data from previous releases where in-place upgrades cannot be supported. These include migrations from:

- an existing Velocity installation to a Watson Explorer Engine installation
- an older Watson Explorer Engine installation to a newer one

Subsequent sections also identify other parts of the Watson Explorer documentation set that describe other migration scenarios, such as the portion of the Results Module documentation that describes the backup and migration process from IO Pro to the Results Module.

Upgrading from Watson Explorer 9.0 to Fix Packs

This section describes the process for upgrading from Watson Explorer v9.0 to fix pack releases of Watson Explorer 9.0, such as v9.0.0.1, v9.0.0.2, and so on.

About this task

The following procedure describes the process for upgrading from Watson Explorer v9.0 to fix pack releases of Watson Explorer 9.0, such as v9.0.0.1, v9.0.0.2, and so on, and for upgrading from a Watson Explorer 9.0.0.x fix pack to a later fix pack.

These instructions discuss using the graphical interface of Installation Manager. The same procedure can also be followed by using the command-line interface if appropriate.

Procedure

1. Download the fix pack from Fix Central and extract the contents of the zip archive for the fix pack.
2. Stop all the services that are associated with your Watson Explorer installation, beginning with Application Builder.

For information about services and the scripts that start and stop the services, see “Integrating Watson Explorer Into the System Startup and Shutdown Process” on page 75.

3. Run Installation Manager as the same user who installed it.

Operating system	Procedure
Linux	From the /opt/IBM/InstallationManager/eclipse directory, enter ./IBMIM. To run in console mode, from the /opt/IBM/InstallationManager/eclipse/tools directory, enter ./imcl -c.
Windows	Click Start > IBM Installation Manager > IBM Installation Manager .

4. Uninstall Watson Explorer 9.0.0:
 - a. In Installation Manager, select **Uninstall**.
 - b. Select **IBM Watson Explorer 9.0.0** to remove Watson Explorer 9.0.0 and click **Next**.
 - c. Click **Uninstall**. A progress bar at the bottom of Installation Manager alerts you when the uninstall process is completed.
 - d. When the uninstall process is completed, a solid green circle with a check mark inside it confirms that Watson Explorer Engine 9.0.0 was uninstalled. Click **Finish**.
5. Add the Watson Explorer package for the fix pack that you are installing to Installation Manager:
 - a. In Installation Manager, click the **File** menu and select **Preferences**.
 - b. Click **Add Repository**.
 - c. In the Add Repository window, click **Browse** and navigate to the ide-offering subdirectory of the directory that was created when you extracted the contents of the fix pack.
 - d. Select the `disktag.inf` file, and click **OK** to add the ide-offering directory as a repository. Click **OK** again to close the Add Repository window.
 - e. Click **OK** to save your changes and close the Preferences window.
 - f. Click **Install**, select the check boxes for the new version of Watson Explorer, and click **Next**.
 - g. Review and accept license agreements and then click **Next**.
 - h. Review the shared resources directory and confirm the installation directory, which should be the same as it was for your previous installation. After you review the directories, click **Next**.
 - i. Select the Watson Explorer modules that you want to install. The modules that you uninstalled previously are selected by default. You can install more modules by selecting the check box for those modules. Alternatively, you can clear any check box for a module that you do not want to install.
 - j. Follow the standard installation process for those modules, as described in “Installing Watson Explorer Graphically” on page 15 or “Installing Watson Explorer from the Command-Line” on page 23.
 - k. After you confirm the settings for the modules that you are installing, click **Install**. A progress bar alerts you to the status of your installation.
 - l. When the installation completes, click **Finish**.
6. Start all the installed services, with Application Builder being last. For information about integrating these services into the automatic startup and shutdown processes of your system, see “Integrating Watson Explorer Into the System Startup and Shutdown Process” on page 75.

Note: For Windows systems, when you upgrade to Version 9.0.0.3 from earlier 9.0 versions, the services for Application Builder, Results Module, and ZooKeeper retain their original display names. For example, the default display name of the service for Application Builder for Version 9.0.0.3 is **Watson Explorer AppBuilder Service**, but if you are upgrading from an earlier 9.0 version, the display name remains **IDE AppBuilder Service**.

Upgrading from Watson Explorer 8.2 Installations

Important: Manually upgrading from Watson Explorer Engine 8.2-x to Watson Explorer 9.0 is necessary on all platforms because you cannot install Watson Explorer Engine 9.0 into a directory in which Watson Explorer Engine 8.2-X was previously installed., even after uninstalling Watson Explorer Engine 8.2-X.

The process of upgrading from Watson Explorer Engine 8.2-X to Watson Explorer Engine 9.0 is the following:

1. Stop all services related to Watson Explorer. If you installed Application Builder, you must stop Application Builder before stopping ZooKeeper.
2. Uninstall Watson Explorer 8.2-X.
3. Prevent installation manager from upgrading, as discussed in “Using the Right Version of Installation Manager” on page 8. Installing the new version of Watson Explorer will upgrade Installation manager to the version that is used by the Watson Explorer 9.0 installation process.
4. Install Watson Explorer 9.0.

Upgrading from existing Watson Explorer 8.2-X installations to Watson Explorer 9.0 requires copying files and directories from your Watson Explorer 8.2-X installation to the new Watson Explorer 9.0 installation. This also enables you to continue using an existing Watson Explorer 8.2-X installation until you have verified that all of the information from that installation has been correctly migrated to the new Watson Explorer 9.0 installation.

To copy mandatory files and directories from a Watson Explorer Engine 8.2-X installation to a new Watson Explorer Engine 9.0 installation, do the following:

1. Install and initialize Watson Explorer Engine 9.0.
2. Use the **velocity-shutdown** (Linux) or (**velocity-shutdown.exe** (Windows) command in the Watson Explorer Engine installation to shut down that instance of Watson Explorer Engine so that you can copy files from your old installation without any possibility of their being used by the new Watson Explorer Engine installation while it is running.
3. If you have installed and are running Results Module, Application Builder, or installed a local copy of the Watson Explorer Knowledge Center, stop the services associated with each of those applications. If you have installed and are running Application Builder, stop the service associated with ZooKeeper once the Application Builder service has stopped.
4. Recursively copy the following files and directories from the Watson Explorer Engine 8.2-X installation location to the same locations in the new Watson Explorer Engine installation:
 - data/collections
 - data/reporting
 - data/repository.xml
 - data/repository.overridden.xml
 - data/search-collections
 - data/system-reporting
 - data/users.xml

Note: The preceding examples use the Linux file/directory name conventions. If you are copying these files from a command prompt or PowerShell on Windows, please use the Windows directory separator, a back slash ('\').

5. Delete the `data/repository.xml.index` and `data/users.xml.index` files in the new Watson Explorer Engine installation.
6. Start Watson Explorer Engine, access it via your web server, and verify that collections are present and searchable.
7. Access the Watson Explorer Engine administration tool in your browser and verify that the list of users in the previous Data Explorer or Velocity installation is still available in the new installation, and that user rights have also been correctly propagated.

If the new Watson Explorer Engine installation works correctly, make sure that any integration that you had done between Watson Explorer Engine and the startup and shutdown processes on your system are updated to reflect the new Watson Explorer Engine installation. Make sure that the old Data Explorer or Velocity installation is no longer being automatically started, shut down, and so on.

Note: If you were also using Application Builder (and were therefore running ZooKeeper) as part of your old Watson Explorer installation, you will also need to copy ZooKeeper's data and database files to the new installation. Contact product support for more information on copying this data.

Finally:

1. If you have installed Application Builder, start the service associated with the ZooKeeper application before starting Application Builder.
2. If you have installed and are running Application Builder, start the service associated with that application once the ZooKeeper application has started successfully.

Important: As mentioned previously, if you were also running ZooKeeper as part of your old Watson Explorer installation, you must migrate your old ZooKeeper data to the new ZooKeeper server before starting Application Builder.

3. If you have installed and are running Results Module, start the service associated with that application.
4. If you installed a local copy of the Watson Explorer Knowledge Center, start the service associated with that application.

Upgrading from Velocity to Watson Explorer Engine

Upgrading from existing Velocity installations to Watson Explorer Engine requires copying files and directories from your existing Velocity installation to the new Watson Explorer Engine installation. This enables you to continue using an existing Velocity installation until you have verified that all of the information from that installation has been correctly migrated to the new Watson Explorer Engine installation.

To copy mandatory files and directories from an existing Velocity installation to a new Watson Explorer Engine installation, do the following:

1. Install and initialize Watson Explorer Engine
2. Use the **velocity-shutdown** command in the Watson Explorer Engine installation to shut down that instance of Watson Explorer Engine so that you can copy files from your old Velocity installation without any possibility of their being used by the new Watson Explorer Engine installation while it is running.

3. If necessary, follow the appropriate update procedures that are necessary to upgrade your Velocity installation to Velocity 8.1-6. Any upgrade procedures for each incremental release between the version of Velocity that you are running and Velocity 8.1-6 are explained in the release notes for each of those releases.
4. Use the **velocity-shutdown** command to shut down your existing instance of Velocity so that you can copy files without any possibility of their being updated.
5. Recursively copy the following files and directories from the old Velocity installation location to the same locations in the new Watson Explorer Engine installation:
 - data/collections
 - data/reporting
 - data/repository.xml
 - data/repository.overridden.xml
 - data/search-collections
 - data/system-reporting
 - data/users.xml
6. Delete the file data/repository.xml.index in the new Watson Explorer Engine installation.
7. Start Watson Explorer Engine, access it via your web server, and verify that collections are present and searchable
8. Access the Watson Explorer Engine administration tool in the and verify that the list of users in the previous Velocity installation is still available in the new installation, and that user rights have also been correctly propagated.

If the new Watson Explorer Engine installation works correctly, make sure that any integration that you had done with Velocity the startup and shutdown processes on your system are updated to reflect the new Watson Explorer Engine installation. make sure that the old Velocity installation is no longer being automatically started, shut down, and so on.

Upgrading from IO Pro to Watson Explorer Results Module

For information about upgrading from IO Pro to Watson Explorer Results Module, see the section of the **Results Module Configuration Guide** in the Knowledge Center, entitled **Upgrading from IO Pro**.

System Administration

Like any enterprise software package, Watson Explorer and its components need to be integrated into your IT group's policies and procedures. Watson Explorer enables you to create and integrate 360-degree information applications that are corporate and business assets and which therefore typically require high availability. The applications must start and stop automatically as part of the boot and shutdown processes on designated systems, the configuration information and (often) data for these applications must be integrated into standard procedures such as backups, the data that is created and used by these applications must be protected from automated anti-virus and anti-malware scanning, and so on. The following sections discuss how system administrators and integrators can integrate Watson Explorer and its components into their application and enterprise IT procedures.

Identifying Installed Watson Explorer Software Versions

When contacting product support, you may need to know the versions of Watson Explorer components that are installed on your system. Watson Explorer provides multiple ways of identifying the versions of the Watson Explorer software and its components that are installed on a given system.

Every Watson Explorer installation includes a text file, located in the root directory of your Watson Explorer software installation, that provides a list of the versions of the Watson Explorer software that were located on the media from which Watson Explorer and related software modules were installed. This file, named `MANIFEST_VERSION.txt`, contains a list of available components and their version numbers. For example, the version of this file that was installed as part of a Watson Explorer 9.0.0.1 fix pack installation is named `MANIFEST_9.0.0.1.txt`, is located in the `/opt/ibm/IDE` directory by default, and contains the following information:

```
ZooKeeper: 3.4.5
WLP: 8.5.5.0
JDK: 7.0.5.0
Engine: 9.0.1.rc10
AppBuilder: 3.0.1.244
ResultsModule: 1.2.1.110
Alfresco (connector): 3.0.1.1
Cba-utils (connector): 1.0.2
Connector-bigdb_cdc (connector): 3.0.1
Connector-confluence (connector): 4.1.0
Connector-exchange (connector): 3.1.1
Connector-lotus (connector): 3.3.4
Connector-mdm-se (connector): 3.1.1
Connector-sap-km (connector): 3.0.2
Connector-wcm-seedlist (connector): 3.1.2
Connector-webcenter_content (connector): 3.0.3
Documentum (connector): 3.1.3
Io-sharepoint (connector): 3.0.5
Sharepoint (connector): 3.0.2
```

Note: The manifest file is not updated if you subsequently install updated connectors that are provided as fix packs.

If you have installed Watson Explorer Engine, you can also use the **vivisimo-versions** command to identify the versions of the Watson Explorer Engine components that are present in the current installation. This command, located in the `Engine/bin` directory of your Watson Explorer installation, produces output like the following:

Software release:

```
Copyright 2000-2013 Vivisimo, Inc. All rights reserved
Compiled Dec 2 2013 at 17:25:13
Release version: 9.0-1
```

Component versions:

```
crawler library: 9.0-1 a85b46f2b621e3149380f5915a6afd62
converter:      9.0-1 4376ff54e2e743177025e6e07495d596
search-engine: 9.0-1 551bd68fd6aafb7ba2c8fdd8b2053bc
libmisc:       9.0-1 0510888f6fa3a81788338b3d312efa15
libvivisimo:   9.0-1 a0db97506405914af291011a4ae001c0
```

```
SHA1:          fbe2bb2c42db405f9954ad7ea308f9c778b38f37
```

Environment Variable Settings for Watson Explorer Engine

This section discusses the cases in which you may need to modify system environment variables in order to successfully execute Watson Explorer Engine applications from command-line or scripted environments on the system(s) where you have installed Watson Explorer Engine.

Configuring Global Options for Watson Explorer Engine

Watson Explorer Engine uses a number of option settings to identify the location of various filesystem resources and, in some cases, to trigger the use of those resources. Many Watson Explorer Engine options, such as the location of temporary directories, are always used, while others, such as the location and name of debugging files, are only used if the **debug** option is used to activate the generation of debugging output. Most of the options used by Watson Explorer Engine are traditionally set in the Watson Explorer Engine administration tool, but some must be set before starting that tool. All of these options can therefore be specified in Watson Explorer Engine's `vivisimo.conf` configuration file regardless of whether the Watson Explorer Engine administration tool is being used.

The default contents of a `vivisimo.conf` file in a Microsoft Windows Watson Explorer Engine installation are the following:

```
<?xml version="1.0" encoding="UTF-8" ?>
  <options>
    <option name="install-dir">C:\Program Files\IBM\IDE\Engine</option>
    <option name="base-url">/vivisimo/</option>
  </options>
```

The default contents of a `vivisimo.conf` file in a Linux Watson Explorer Engine installation are the following:

```
<?xml version="1.0" encoding="UTF-8" ?>
  <options>
    <option name="install-dir">/opt/ibm/IDE/Engine</option>
    <option name="base-url">/vivisimo/</option>
  </options>
```

To set any of the options that are discussed in this section, simply add an appropriate **option** entry to the `vivisimo.conf` file for your Watson Explorer Engine installation. Any directories that you modify from their installation defaults must be writable by the web server with which Watson Explorer Engine is integrated.

warning: The `vivisimo.conf` file should only be modified when Watson Explorer Engine is not running. If you need to modify the contents of this file, IBM recommends that you modify this file before running Watson Explorer Engine for the first time.

For example, the following `vivisimo.conf` file sets the directory used to store debugging output (**debug-dir**) to `C:\debug` (which must already exist - Watson Explorer Engine will not create this directory):

```
<?xml version="1.0" encoding="UTF-8" ?>
  <options>
    <option name="install-dir">C:\Program Files\IBM\IDE\Engine</option>
    <option name="base-url">/vivisimo/</option>
    <option name="debug-dir">C:\debug</option>
  </options>
```

Note: When setting networked file or directory locations for Watson Explorer Engine related storage on Microsoft Windows systems, do not use mapped drives

(in other words, network locations that are associated with a drive letter) because they may be user or system-specific. Watson Explorer Engine only supports networked file or directory locations for Watson Explorer Engine-related storage that are specified in Universal Naming Convention (UNC) notation.

Using Shared Libraries

The binary applications that are provided with Watson Explorer Engine are installed in the **bin** subdirectory of a Watson Explorer Engine installation directory. On Linux systems, the shared libraries used by those applications are located in the **lib** subdirectory of an installation directory. Watson Explorer Engine includes wrapper scripts that automatically add the directory where shared libraries are stored to your working **LD_LIBRARY_PATH** for the primary Watson Explorer Engine executables that system administrators may want to execute, namely:

- admin-cmd
- checksum
- chico
- snmp-velocity
- velocity-cmd
- velocity-config
- velocity-shutdown
- velocity-startup
- velocity-webserver
- vivisimo-versions

Executing any of these applications from the **bin** subdirectory of a Watson Explorer Engine installation will first set the **LD_LIBRARY_PATH** to include the **lib** subdirectory of that same installation, and will then execute the corresponding binary from the **bin/real** subdirectory.

warning: IBM recommends that you do not directly execute any of the other applications in the **bin** subdirectory of a Watson Explorer Engine installation. These applications are only provided for use by Watson Explorer Engine and are only supported in that context.

Locating Executables

The location of executables on Microsoft Windows and Linux systems can be configured in various ways, but is most easily specified by supplying the full pathname of any Watson Explorer Engine command that you want to execute, either from the command-line or in any graphical application where you are registering a Watson Explorer Engine application (as discussed in “Startup/Shutdown Integration on Microsoft Windows Systems” on page 77).

We recommend that you always identify Watson Explorer Engine applications by their full pathname. However, if you cannot do so for some reason or want to be able to execute them by simply specifying the name of an executable, the following sections explain how to add the directory containing your Watson Explorer Engine executables to the **PATH** environment variable that specifies the search path for executables on Linux and Microsoft Windows systems.

Locating Executables on Linux Systems: On Linux systems, the **PATH** environment variable is most commonly set in the configuration files for each user's shell. Linux systems typically use the **Bourne-Again Shell** shell, **Bash** (/bin/bash) for most users. The system-wide configuration file for the Bash shell is the file **/etc/bashrc**, but you rarely add the **bin** directory for third-party

applications such as Watson Explorer Engine to the settings in this file, because all users would then be able to attempt to execute the administrative Watson Explorer Engine commands. Instead, you would add Watson Explorer Engine's **bin** directory to the per-user configuration file for the Bash shell for selected users, such as the root user. The per-user configuration file for bash is the file **.bashrc**, located in each user's home directory.

To add Watson Explorer Engine's **bin** directory to the search path for executables that is used by the root user, add the following entry to the end of the file `/root/.bashrc` as the root user or via the **sudo** command:

```
export PATH=${PATH}:/path/to/data-explorer-engine/bin:
```

The new **PATH** setting will take effect the next time that the root user logs in, and will also be in effect in any shells or command scripts that are executed as the root user.

Locating Executables on Windows Systems:

About this task

Due to its focus on a graphical execution environment, the **PATH** environment on Windows systems is set graphically. Unfortunately, different versions of Windows set environment variables in different ways. For example, to add Watson Explorer Engine's **bin** directory to your system's **PATH** variable on a Windows 2003 or later Windows server system, do the following:

Procedure

1. From the desktop, right-click **My Computer** and click **Properties**, or select the **System** Applet in the Control Panel.
2. In the System Properties window, click on the Advanced tab.
3. In the Advanced section, click the Environment Variables button.
4. In the Environment Variables window, highlight the **PATH** variable in the Systems Variable section and click **Edit**.
5. Add the entry for Watson Explorer Engine's **bin** directory to the end of the **PATH** variable, separated from the previous directory by a semicolon, as in the following example:

```
C:\Program Files;C:\Winnt;C:\Winnt\System32;C:\Program Files\IBM\IDE\Engine\bin
```

6. Click **OK** to save your changes. The update environment variable will be available after rebooting your system.

Results

If you are using a different version of Windows than the one shown in this example, see your Windows documentation for detailed information about setting environment variables.

System and Application Configuration for Watson Explorer Engine

Depending on how you use Watson Explorer Engine, you may need to change various system or sub-system parameters to maximize its performance or the performance of other system applications, such as your web server, on which Watson Explorer Engine depends. This section discusses parameters that you should be aware of and, in some cases, may want to modify. Where possible, this section also explains how to modify such parameters for operating systems and relevant enterprise applications.

Increasing the Maximum Number of Open Files

Watson Explorer Engine applications crawl and enqueue large numbers of files, and may therefore need to keep more files open than your system or your web server support by default. Common signs that you may have exceeded the current limits of your system are Watson Explorer Engine log or error messages along the lines of "Too many open files".

Maximum Open File Settings on Linux Systems: Watson Explorer Engine uses the standard Linux `select()` call to manage file descriptors, which is effectively limited to handling up to 1024 file descriptors. On systems where Watson Explorer Engine is installed, you should not attempt to increase the number of file descriptors that are available.

Microsoft Windows Internet Information Services (IIS): On Microsoft Windows systems, the maximum number of open files is a system configuration value that cannot be changed. Windows applications that use the C runtime library can open a maximum of 10000 files simultaneously. Windows applications that use native Win32 API calls can open up to 16384 files simultaneously. Watson Explorer Engine uses the C runtime library, and is therefore limited to a maximum of 10000 open files.

If you are seeing "Too many open files" messages on a Windows system where Watson Explorer Engine is running, the simplest solution is to minimize the number of other processes that are running on that system. Configuring your system so that its primary job is to run Watson Explorer Engine and either Watson Explorer Engine's internal web server or Microsoft IIS will reduce the chance that other applications will use file handles that are needed by Watson Explorer Engine.

Configuring Watson Explorer Engine's Embedded Web Server

Watson Explorer Engine includes an embedded web server that can be used to isolate Watson Explorer Engine from any existing web server that is running on your system, or to provide a dedicated web server for use with Watson Explorer Engine API applications. How this web server runs on your system and any security that it uses are configured using the `velocity-config` command, which is explained in the next section.

Important: If you are using Watson Explorer Engine on a Microsoft Windows system, you should not attempt to use both Microsoft's IIS web server and Watson Explorer Engine's embedded web server to interact with the Watson Explorer Engine administration tool. These web servers typically execute as different users and groups, start processes using different user/group identities, and require different file permission and ownership settings on your Watson Explorer Engine installation. Use either one of these web servers exclusively when using the Watson Explorer Engine administration tool.

Any of the Watson Explorer Engine administration commands discussed in this section or in "System Administration" on page 56 must be executed with at least the system privileges that you used when installing Watson Explorer Engine. In other words, if you installed Watson Explorer Engine as the root user on a Linux system or as Administrator on a Microsoft Windows system, you must execute the Watson Explorer Engine administration commands as that user or as another user with equivalent or superior privileges.

See "Web Server User Identity and File Permissions" on page 68 for information about starting the embedded web server with administrative privileges but

executing the web server process as another user. If you encounter an error when executing any of the **velocity-config** command, see “Command-Line Watson Explorer Engine Administration Commands” on page 68 for general information about resolving errors related to Watson Explorer Engine’s administrative commands.

To start and access the embedded web server:

1. Use the **velocity-config** command (Linux) or **velocity-config.exe** command (Microsoft Windows) to:

- a. Enable the embedded web server, using the following command:

```
velocity-config webservice/is-enabled=True
```

- b. Configure the user and group that the embedded web server should run as, using commands like the following:

```
velocity-config global/user=www-data  
velocity-config global/group=www-data
```

On a Linux system, you must configure the user and group that the embedded web server should run as. If you have installed Watson Explorer in non-admin mode, you would use the user and group information that you collected in “Collecting User/Group Information for Non-Admin Linux Installations” on page 10. Otherwise, use the user and group information that you collected in “Collecting User/Group Information for Privileged Linux Installations” on page 10

Setting the user and group configuration options is optional on a Microsoft Windows system, and depends on the user as whom you installed Watson Explorer Engine. See “Web Server User Identity and File Permissions” on page 68 for detailed information about the user, file, and directory permissions that are required to create and modify Watson Explorer Engine data.

You may also want to set other, optional configuration parameters. See “velocity-config Command” on page 63 for detailed information about available configuration options and about using the command in general.

2. Do one of the following to actually start the web server:
 - (Recommended) Use the **velocity-startup** command (Linux) or **velocity-startup.exe** command (Microsoft Windows) to start the web server, enable access to Watson Explorer Engine, and optionally start other Watson Explorer Engine related processes. (See “velocity-startup” on page 69 for detailed information about using this command.)
 - Use the **velocity-webservice** command (Linux) or **velocity-webservice.exe** command (Microsoft Windows) to manually start the web server, which is explained in “velocity-webservice” on page 73. This command will start the web server, but you may still not be able to access Watson Explorer Engine, depending on how it was last shut down on your system.

Once you have started Watson Explorer Engine’s embedded webservice, you can complete your Watson Explorer Engine installation by accessing the Watson Explorer Engine administration tool through the embedded web server. If you do not know how to contact Watson Explorer Engine’s embedded web server, you can use the **info** option of the **velocity-webservice** command (Linux) or **velocity-webservice.exe** command (Microsoft Windows) to show a URL that you can use to contact the web server.

For information about integrating the **velocity-startup** command into your system’s startup and shutdown processes, see “Command-Line Watson Explorer Engine Administration Commands” on page 68.

Customizing Embedded Web Server Execution

Watson Explorer Engine's embedded web server provides a web server that Watson Explorer Engine and Watson Explorer Engine API applications can use without affecting any other web server that may already be running on a given system, if no web server is installed on that system, or if you have installed Watson Explorer in non-admin mode. If you are configuring the Watson Explorer Engine embedded web server on a Linux system where no other web server has previously been installed or where you want the embedded web server to be completely independent of any other web server, you should create a unique user and group that the embedded web server can run as.

To create a group and user for running Watson Explorer Engine's embedded web server on a Linux system, execute commands like the following as the root user:

```
groupadd -g 750 deengine
useradd -u 750 -g deengine -s /bin/bash -d /tmp
```

These example commands will create a **deengine** group and then create an associated **deengine** user.

Note: You should replace the number 750, which is, respectively, the group ID (GID) and user ID (UID) in these command examples, with a numeric GID and UID that is keeping with your system administration procedures.

After creating this user and group, you can now use the **velocity-config** command to set the user and group that the web server should run as, using commands like the following:

```
velocity-config global/user=deengine
velocity-config global/group=deengine
```

You can now configure other aspects of the embedded web server, such as the port that it will listen on, whether it is enabled, any security certificate and key that is associated with it, and so on. These aspects of configuring the embedded web server are discussed in detail in “Configuring Watson Explorer Engine's Embedded Web Server” on page 61.

Note: If you will be running Watson Explorer Engine's embedded web server as a specified user and group, that user and group must own the Watson Explorer Engine installation directory. You can recursively change the ownership of this directory to that user and group by changing directory to your Watson Explorer Engine installation directory and executing the following command, where **USER** and **GROUP** are, respectively, the user and group that you want to run Watson Explorer Engine as:

```
chown -R USER:GROUP bin cache debug doc jre lib tmp webserver www
```

velocity-config Command

Enables you to configure Watson Explorer Engine's embedded web server, including enabling whether the web server is started automatically as part of the **velocity-startup** command.

Velocity Config Command

velocity-config Enables you to configure the Watson Explorer Engine embedded web server, including enabling whether the web server is started automatically as part of the **velocity-startup** command.

Synopsis

```
velocity-config
  [-h|--help]
  [-I|--installation-dir /path/to/data-explorer-engine/installation]
  [-q|--quiet]
  [--version]
  [name=value]
```

Description

The **velocity-config** command is located in the `bin` subdirectory of a Watson Explorer Engine installation. This command enables you to display and set configuration information for Watson Explorer Engine's embedded web server by specifying *name=value* pairs on the command-line. This configuration information is stored in the file `data/velocity-config.xml`. This file will only exist if you have modified the configuration of the embedded web server. If the file does not exist, the default configuration values (listed later in this section) are used.

Executing the **velocity-config** command with no options displays the configuration options that are available for the embedded web server and their current values. The command and its output will look something like the following (which shows the default values for all of these options):

```
$ velocity-config
global/admin-email          admin@localhost
global/admin-url           http://HOSTNAME-OR-IP/vivisimo/cgi-bin/admin
global/api-url             <unset>
global/group               <unset>
global/password            <unset>
global/user                <unset>
webserver/base-url        /vivisimo/
webserver/host            HOSTNAME-OR-IP
webserver/is-enabled      False
webserver/is-ssl-enabled  False
webserver/port            9080
webserver/service-name    VelocityWebServer
webserver/ssl-certificate <unset>
webserver/ssl-certificate-key <unset>
```

Note: On Microsoft Windows systems, the **velocity-config** command does not display the **global/group** option. On Linux systems, this command does not display the **global/password** and **webserver/service-name** options. See the description of these options later in this section for more information about their use and possible values.

As you can see from this example, the **velocity-config** command enables you to set two classes of configuration options: **global** values, which will affect the web server but which are intended to be used by other Watson Explorer Engine components, and **webserver** values, which only affect the configuration of the web server. To set any of these values, specify the name of the option that you want to set and the value that you want to set it to, separated by an equals sign ('='). For example, to set the port that the webserver uses to port 10025, you would execute the command `velocity-config webserver/port=10025`. This command and its output would look like the following:

```
$ velocity-config webserver/port=10025
Configuration saved...
webserver/port 10025
```

The **velocity-config** command does not validate the values of any parameters that you set. For example, port numbers must be less than 65536, but are not checked when you set them.

Note: The **global** values that you can configure using the **velocity-config** command are currently only used by the Watson Explorer Engine embedded web server.

On Microsoft Windows systems, the **velocity-config** command has the **.exe** file extension.

Options

- **-h|--help:** Displays a usage message and exits without performing any other actions.
- **-I|--installation-dir/path/to/data-explorer-engine/installation:** Enables you to specify the location of the Watson Explorer Engine installation that contains the web server which you want to configure. If no installation directory is specified, the **velocity-config** command assumes that the webserver directory and binaries are located in the Watson Explorer Engine installation directory which is the parent of the directory in which the **velocity-config** application is located.
- **-q|--quiet:** Suppresses normal output and warnings. Errors will still be displayed.
- **--version:** Displays the version of Watson Explorer Engine from which the **velocity-config** command is being executed.

Configuration Items

The following list shows the configuration items for which values can be specified using the **velocity-config** command:

Important: If you are using Watson Explorer Engine's embedded web server, you must restart it after changing any of the configuration items supported by the **velocity-config** command.

- **global/admin-email:** The email address shown by the webserver when an error is displayed. The default value of this configuration option is **admin@localhost**.
- **global/admin-url:** The URL that should be used to contact the Watson Explorer Engine administration tool using the embedded web server. This value cannot be set manually - it is constructed internally using the **webserver/base-url** and **webserver/host** configuration options.
- **global/api-url:** The URL that should be used to execute Watson Explorer Engine startup, shutdown, and API commands. The default value for this configuration option is **<unset>**. When this value is not set:
 - If Watson Explorer Engine's embedded web server is not running, Watson Explorer Engine will attempt to derive this URL from **localhost**, the default port number for the web server, the virtual directory (base-url) that is specified in the Watson Explorer Engine installation's **vivisimo.conf** file, and the **cgi-bin** subdirectory.
 - If Watson Explorer Engine's embedded web server is running, Watson Explorer Engine will use the URL that is displayed as part of the output of the **velocity-webserver info** command. This URL is derived from the hostname for the current system (if available), the port number specified in the **webserver/port** option, the virtual directory (base-url) that is specified in the installation's **vivisimo.conf** file, and the **cgi-bin** subdirectory.

This option is provided for use when Watson Explorer Engine applications are not installed in their default location, and therefore cannot be accessed through the default URL.

- **global/group**: The name of the group that the web server should run as. On Linux systems, you must set both this configuration option and the **global/user** configuration option in order to start Watson Explorer Engine's embedded web server as the root user. This option is not used on Microsoft Windows systems. The default value of this configuration option is **<unset>**, indicating that no value has been set for this option yet. When this value is not set, the webserver will attempt to run as the group that is associated with the user who starts the webserver.
- **global/password**: The password associated with the user that the web server should run as (which was specified by the **global/user** configuration option). The actual text value of the password is shown as a sequence of asterisks for security purposes, and is stored internally in a secure fashion. This option is only used on Microsoft Windows systems. The default value of this configuration option is **<unset>**, indicating that no value has been set for this option yet.
- **global/user**: The name of the user that the embedded web server should run as. On Linux systems, the embedded web server cannot run as the root user, but is typically started by a system script that is executed by the root user. You must therefore set both the **global/user** and **global/group** configuration options in order to identify the user and group identity that the root user will use to start and run the embedded web server. The specified user must be a member of the group specified by the **global/group** configuration option.

Note: Attempting to set **global/user** to the root user will display error messages, as in the following example:

```
# velocity-config global/user=root
ERROR: The value for 'global/user' is invalid:
      Running Watson Explorer Engine as root is insecure and not allowed
      (ID: CLI_INVALID_CONFIG_VALUE_ERROR)

WARNING: No changes were made to the stored configuration.
      (ID: CLI_NO_CONFIG_CHANGES_WARNING)
```

On Microsoft Windows systems, you must set both the **global/user** configuration option and the **global/password** option in order to start Watson Explorer Engine's embedded web server as a specified user. Some Windows installations may require that you specify the fully-qualified user name (*DOMAIN\user*) in order to correctly authenticate as a specified user.

The default value of this configuration option is **<unset>**, indicating that no value has been set for this option yet. When this value is not set on Linux systems, the webserver will attempt to run as the user who starts the webserver.

- **webserver/base-url**: The base part of the admin and query-meta URLs. This configuration option is listed for convenience only, and cannot be set by the **velocity-config** command. It can only be set by editing the *vivisimo.conf* configuration file. The default value of this option is */vivisimo/*. This option **must** begin with a */* symbol, or attempting to access Watson Explorer Engine will result in an error.
- **webserver/host**: The host name or IP address of the system on which the embedded Watson Explorer Engine web server is running.
- **webserver/is-enabled**: Whether the embedded web server is enabled (true) or not (false). If this variable is set to **true** and the embedded web server has not been removed from your Watson Explorer Engine installation, executing the

velocity-startup command will automatically start the embedded web server. The default value of this configuration option is **false**.

- **webserver/is-ssl-enabled:** Whether the embedded web server uses the Secure Socket Layer (SSL), now known as Transport Layer Security (TLS), for secure communication and data integrity verification. If this variable is set to **true**, you must also provide both an SSL certificate and key. The default value of this configuration option is **false**.
- **webserver/port:** The port on which the embedded web server listens for data. The default value of this configuration option is **9080**.

Note: On Microsoft Windows hosts where multiple users are doing development and are therefore running different instances of Watson Explorer Engine, each instance must use a unique **webserver/port** and **webserver/service-name** value. Using the same port number or service name will cause a variety of different Microsoft Windows errors depending upon the type of Watson Explorer Engine operation that you subsequently try to perform.

- **webserver/service-name:** (Microsoft Windows only) The name of the Windows Service that is associated with Watson Explorer Engine. The default value of this option is **VelocityWebServer**, which is a service that is created by Watson Explorer Engine.
- **webserver/ssl-certificate:** The SSL certificate for the embedded web server. To provide a certificate, save it to a file and specify the name of the file, preceded by an '@' sign, as the value for this configuration option. As an example, the following command would set the SSL certificate for an embedded Watson Explorer Engine web server to the contents of the file `corporate.crt`:

```
velocity-config webserver/ssl-certificate=@corporate.crt
```

You can specify any file name as long as it contains the SSL certificate that you want to use. The default value of this configuration option is **<unset>**, indicating that no value has been set for this option yet.
- **webserver/ssl-certificate-key:** the key for the SSL certificate associated with the embedded web server. To provide a key, save it to a file and specify the name of the file, preceded by an '@' sign, as the value for this configuration option. As an example, the following command would set the SSL certificate key for an embedded Watson Explorer Engine web server to the contents of the file `corporate.key`:

```
velocity-config webserver/ssl-certificate-key=@corporate.key
```

You can specify any file name as long as it contains a valid key for the SSL certificate that you are using. The default value of this configuration option is **<unset>**, indicating that no value has been set for this option yet.

Messages

See the **Watson Explorer Engine Message Reference Manual** for detailed information about any other messages that you may receive from the **velocity-config** command.

Exit Status

Exits with status **0** when the command is successful.

Note: On Linux systems, exit codes above 127 are returned by the shell. Please consult the documentation for the shell that you are using for more information.

Web Server User Identity and File Permissions

As discussed in “velocity-config Command” on page 63, the **velocity-config** command enables you to enable the embedded web server, identify the user that you want to run the embedded webserver as, specify a password for that user, and so on. Enabling the web server means that it will be automatically started by the **velocity-startup** command, which is typically integrated into your system's startup and shutdown process on systems where Watson Explorer Engine is used. These system startup and shutdown processes typically run as the root user on Linux systems and as an Administrator on Microsoft Windows systems. Configuring the embedded web server to run as a specified user enables the web server to be started by a privileged user but execute as another user. This is the standard execution model for popular web servers such as the Apache web server.

When following this execution model, the user that you configure the embedded webserver to run as must have sufficient privileges to access the files and directories in your Watson Explorer Engine installation. After configuring the embedded webserver, you must therefore make the specified user the owner of the files and directories in your installation or add that user to file and directory ACLs so that they can write to and create files in your installation directory.

Command-Line Watson Explorer Engine Administration Commands

This section describes some of the Watson Explorer Engine functions that can be performed from the command prompt. Commands discussed in this section include **velocity-startup**, **velocity-shutdown**, **velocity-webserver**, and **velocity-cmd**. A related command, **velocity-config** command, discussed in the context of the “Configuring Watson Explorer Engine's Embedded Web Server” on page 61 documentation, enables you to set the configuration variables that provide mandatory information for Watson Explorer Engine components such as the embedded web server and modules such as **Watson Explorer Results Module**.

Important: When using any of Watson Explorer Engine's command-line administrative commands, encountering an error such as (ID: API_INVALID_SU_TOKEN) The current session could not be authenticated with a su-token often indicates that the Watson Explorer Engine installation was initialized using a configuration **api-url** whose hostname or IP address can no longer be resolved from the system on which Watson Explorer Engine is installed. You can eliminate this error message by using the **velocity-config** command, as described in “velocity-config Command” on page 63.

Watson Explorer Engine provides two commands that can be used from any command prompt to start or stop Watson Explorer Engine on a given system, and which are designed to be easily integrated into a system's startup and shutdown procedures to guarantee the availability of Watson Explorer Engine or other applications that use Watson Explorer Engine. These applications are the **velocity-startup** and **velocity-shutdown** commands, which are located in the **bin** subdirectory of your Watson Explorer Engine installation directory.

Note: The **velocity-startup** command will always start the embedded Watson Explorer Engine web server if the web server is installed and enabled. For more information about configuring and enabling this embedded web server, see “Configuring Watson Explorer Engine's Embedded Web Server” on page 61. You can also use the “**velocity-webserver**” on page 73 command to start, stop, and restart the embedded web server manually.

On Linux and Microsoft Windows systems, you can execute the **velocity-startup** and **velocity-shutdown** commands as any user that has write access to the `data/su-tokens` subdirectory of the Watson Explorer Engine installation that you are starting or stopping. This is typically the user that you installed the Watson Explorer Engine software as. The root user on Linux systems and the Administrator user on Microsoft Windows systems should always have write access to this directory.

Note: When you shut down Watson Explorer Engine, it records the reason for that shut down in files named `maintenance-mode` that are located in the top level and `bin` directories of `/path/to/data/explorer/engine/installation`. The presence of these files indicates that Watson Explorer Engine has been shut down cleanly.

The next sections provide complete reference information for Watson Explorer Engine's command-line startup and shutdown commands. For information about integrating these commands into your system's startup process, see “Integrating Watson Explorer Into the System Startup and Shutdown Process” on page 75. A subsequent section also discusses the **velocity-webserver** command, which can be used to manually start, stop, and restart Watson Explorer Engine's embedded web server. For information about configuring the embedded web server, see “Configuring Watson Explorer Engine's Embedded Web Server” on page 61.

Watson Explorer Engine provides single commands that enable you to start and stop all of its processes in an orderly fashion. For information about the command used to start Watson Explorer Engine, see “`velocity-startup`.” For information about the command used to shut down Watson Explorer Engine, see “`velocity-shutdown`” on page 71. For information about integrating these commands into your system's startup process, see “Integrating Watson Explorer Into the System Startup and Shutdown Process” on page 75.

Tip: The “`velocity-shutdown`” on page 71 ordinarily signals all running Watson Explorer Engine processes to terminate cleanly, which enables these processes to checkpoint in-progress work. If you need to terminate all Watson Explorer Engine processes immediately, you can use the “`velocity-shutdown`” on page 71's **--force** option. However, using the **--force** option may require additional time when restarting Watson Explorer Engine because the Watson Explorer Engine processes may need to recover and resubmit in-progress work.

velocity-startup

Removes Watson Explorer Engine from maintenance mode, allowing other Watson Explorer Engine processes to start. Additional command-line arguments can be supplied to start the Watson Explorer Engine query service and scheduler.

Synopsis

```
velocity-startup [-h|--help]
                 [-I|--installation-dir /path/to/data-explorer-engine/installation]
                 [-q|--quiet]
                 [--start-query-service]
                 [--start-scheduler]
                 [--version]
                 [-y|--yes]
```

Description

The **velocity-startup** command is located in the `bin` subdirectory of a Watson Explorer Engine installation. (On Microsoft Windows systems, this command has the `.exe` file extension.) Executing this command with the right options takes a

Watson Explorer Engine installation out of maintenance mode, enabling Watson Explorer Engine processes to be started, but does not actually start any Watson Explorer Engine processes.

NOTE: The **velocity-shutdown** command requires the version of the **curl** application that is provided with a Watson Explorer Engine installation.

Executing the **velocity-startup** command with the **--start-scheduler** option starts the Watson Explorer Engine scheduler with a "boot" message, causing the scheduler to start all Watson Explorer Engine services that have been registered to start at boot time. Schedulable services include repository synchronization, building a dictionary, building and emailing a report, running alerts, starting or stopping a search collection crawl, starting or stopping the query service, or testing a specified source.

Executing the **velocity-startup** command with the **--start-query-service** option starts the Watson Explorer Engine query service, which enables a Watson Explorer Engine installation to respond to queries.

Note: When starting the embedded Watson Explorer Engine web server for the first time on Microsoft Windows systems that are running the Windows firewall or which support User Account Control (UAC) warnings, your system will display a dialog requesting that you unblock the program in your firewall or otherwise grant the application access to your network. You must approve whichever of these access requests you see in order for applications to be able to use this web server.

Options

- **-h|--help:** Displays a usage message and exits without performing any other actions.
- **-I|--installation-dir/path/to/data-explorer-engine/installation:** Enables you to specify the location of the Watson Explorer Engine installation that you want to start. If no installation directory is specified, the **velocity-startup** command assumes that the Watson Explorer Engine installation directory is the parent of the directory in which the **velocity-startup** application is located.
- **-q|--quiet:** Suppresses normal output and warnings. Errors will still be displayed. Specifying this option implies the **--yes** option. On Windows systems, this option also suppresses the "Press return to continue" message that is shown at the end of each Watson Explorer Engine administrative command.
- **--start-query-service:** Starts the Watson Explorer Engine query service, enabling the copy of Watson Explorer Engine that is associated with this instance of the **velocity-startup** command (or the version in the directory specified using the **-I** option to respond to incoming queries.
- **--start-scheduler:** Starts the Watson Explorer Engine scheduler with a "boot" message, causing the scheduler to start all Watson Explorer Engine services that have been registered to start at boot time.
- **--version:** Displays program version information and exits without performing any other actions.
- **-y|--yes:** Proceeds without prompting for confirmation

Messages

The **velocity-startup** command displays the message "Velocity has been started" when it is successful. Depending on the command-line options that you specify, you may be prompted for other information or see additional messages, such as the following:

- "Is it okay to start Velocity? (y/N)": This confirmation message is displayed if you do not specify the **-y** or **--yes** options. The default response if you simply press return is 'N' (no), which causes the **velocity-startup** command to exit without removing Watson Explorer Engine from maintenance mode.
- "Velocity isn't currently shutdown. Should I continue anyway? (y/N)": Watson Explorer Engine uses files named maintenance-mode that are located in the top level and bin directories of */path/to/data-explorer-engine/installation* to identify the fact that it has been shut down. This message is displayed if you execute the **velocity-startup** command when these files do not exist.
- "Press Return to continue": (Microsoft Windows only) Shown after the **velocity-startup** completes the startup process but before it exits, in order to enable you to see any messages that may have been generated in the command window during the startup process.

See the **Watson Explorer Engine Message Reference Manual** for detailed information about any other messages that you may receive from the **velocity-startup** command.

Exit Status

Exits with status **0** when the command is successful.

Note: On Linux systems, exit codes above 127 are returned by the shell. Please consult the documentation for the shell that you are using for more information.

velocity-shutdown

Terminates Watson Explorer Engine and all related services in an orderly fashion

Synopsis

```
velocity-shutdown [--force]
  [-h|--help]
  [-I|--installation-dir /path/to/data-explorer-engine/installation]
  [-q|--quiet]
  [--reason "Reason for shutdown"]
  [-y|--yes]
```

Description

The **velocity-shutdown** command is located in the bin subdirectory of a Watson Explorer Engine installation. (On Microsoft Windows systems, this command has the **.exe** file extension.) Executing this command with the **-y** option or by answering 'y' to all of its prompts initiates the Watson Explorer Engine shutdown process, which cleanly terminates all Watson Explorer Engine, prevents these services from restarting, and prevents any other Watson Explorer Engine services from being started.

NOTE: The **velocity-shutdown** command requires the version of the **curl** application that is provided with a Watson Explorer Engine installation.

Options

- **--force**: Terminates all Watson Explorer Engine processes immediately, without waiting for those processes to terminate cleanly. Using this option can cause delays when restarting Watson Explorer Engine if crawls were in progress when Watson Explorer Engine was shut down. You should specify this option if you are using the Collection Broker.
- **-h|--help**: Displays a usage message and exits without attempting to shut down Watson Explorer Engine.
- **-I|--installation-dir/path/to/data-explorer-engine/installation**: Enables you to specify the location of the Watson Explorer Engine installation that you want to shut down. If no installation directory is specified, the **velocity-shutdown** command assumes that the Watson Explorer Engine installation directory is the parent of the directory in which the **velocity-shutdown** command is located.
- **-q|--quiet**: Suppresses normal output and warnings. Errors will still be displayed. Specifying this option implies the **--yes** option. On Windows systems, this option also suppresses the "Press return to continue" message that is shown at the end of each Watson Explorer Engine administrative command.
- **--reason "Reason for shutdown"**: provides *Reason for shutdown* as the reason that Watson Explorer Engine was shut down, storing this information in files named *maintenance-mode* that are located in the top level and *bin* directories of */path/to/data-explorer-engine/installation*. If you do not specify a reason, the default reason "System Administrator requested shutdown" is used.
- **--version**: Displays program version information and exits without performing any other actions.
- **-y|--yes**: Proceeds without prompting for confirmation

Messages

The **velocity-shutdown** command displays the message "Velocity has been shutdown" when a shutdown is successful. Depending on the command-line options that you specify, you may be prompted for other information or see additional messages, such as the following:

- "Is it okay to shutdown Velocity?": This message is displayed if you do not specify the **-y** or **--yes** options. The default response if you simply press return is 'N' (no), which causes the **velocity-shutdown** command to exit without attempting to stop Watson Explorer Engine.
- "Velocity is already shutdown. Should I continue anyway?": Watson Explorer Engine records the reason why it has been shut down in files named *maintenance-mode* that are located in the top level and *bin* directories of */path/to/data-explorer-engine/installation*. This message is displayed if you execute the **velocity-shutdown** command when these files already exist.
- "Press Return to continue": (Microsoft Windows only) Shown after the **velocity-shutdown** completes the shutdown process but before it exits, in order to enable you to see any messages that may have been generated in the command window during the shutdown process.

See the **Watson Explorer Engine Message Reference Manual** for detailed information about any other messages that you may receive from the **velocity-shutdown** command.

Exit Status

Exits with status **0** when the command is successful.

Note: On Linux systems, exit codes above 127 are returned by the shell. Please consult the documentation for the shell that you are using for more information.

velocity-webserver

Manages the run state of the embedded web server that is installed as a part of Watson Explorer Engine.

Synopsis

```
velocity-webserver [-h|--help]
                  [-I|--installation-dir /path/to/data-explorer-engine/installation]
                  [-q|--quiet]
                  [--version]
                  [start|stop|restart|graceful|configtest|uninstall|install|info]
```

Description

The **velocity-webserver** command is located in the `bin` subdirectory of a Watson Explorer Engine installation. (On Microsoft Windows systems, this command has the `.exe` file extension.) Executing this command with one of the **start** | **stop** | **restart** | **graceful** changes the execution state of the Watson Explorer Engine web server in the indicated fashion. Executing this command with the **configtest** option enables you to verify that the web server's configuration file is syntactically correct. Executing this command with the **info** options displays information about the status of the web server and the URL at which it can be contacted. The **graceful** option is only available on Linux systems, and restarts the Watson Explorer Engine embedded web server without terminating any active connections. The **uninstall** and **install** options are only available on Microsoft Windows systems, and respectively uninstall or install Watson Explorer Engine's embedded web server as a Windows service.

Note: When starting the embedded Watson Explorer Engine web server for the first time on Microsoft Windows systems that are running the Windows firewall or which support User Account Control (UAC) warnings, your system will display a dialog requesting that you unblock the program in your firewall or otherwise grant the application access to your network. You must approve whichever of these access requests you see in order for applications to use this web server.

Options

- **-h|--help:** Displays a usage message and exits without performing any other actions.
- **-I|--installation-dir/path/to/data-explorer-engine/installation:** Enables you to specify the location of the Watson Explorer Engine installation that contains the web server which you want to configure. If no installation directory is specified, the **velocity-webserver** command assumes that the webserver directory and binaries are located in the Watson Explorer Engine installation directory which is the parent of the directory in which the **velocity-webserver** application is located.
- **-q|--quiet:** Suppresses normal output and warnings. Errors will still be displayed. On Windows systems, this option also suppresses the "Press return to continue" message that is shown at the end of each Watson Explorer Engine administrative command.
- **--version:** Displays the version of Watson Explorer Engine from which the **velocity-webserver** command is being executed.
- **configtest:** Checks the syntax of the web server configuration files (`INSTALLDIR/webserver/webserver.conf`) and any customized configuration

files (`webservice.something.conf`) that you may have created in the data subdirectory of your Watson Explorer Engine installation directory. If an error is present in any of these files, the command displays detailed information about the syntax error. If no syntax errors are present, the command displays the message "Syntax OK".

- **graceful**: (Linux only) Gracefully restarts the embedded web server by sending it a signal, which does not abort or close any connections that are currently open to the web server (including connections from the web server to its log files). If the web server process (`httpd`) is not running, it is started. This command automatically checks the configuration files via `configtest` before initiating the restart to make sure that the embedded web server will be able to restart with those configuration files.

Note: The **graceful** option is not supported by the **velocity-webservice** command on Microsoft Windows systems.

- **info**: Displays information about the status of the webservice and the URL to use to contact the Watson Explorer Engine administration tool through the webservice. This command will display one of the following messages, depending upon its configuration and execution state:

1. If the webservice has not been enabled, it cannot be running, so you will see a message like the following:

The webservice is not enabled.

You can enable the embedded webservice using the **velocity-config**, and then start it using the **velocity-webservice start** command. Once enabled, the web server will be started automatically by the **velocity** command.

2. If the webservice has been enabled, but is not running, you will see a message like the following:

The webservice is enabled and not running.

The url for the Velocity administrative interface is:
`http://localhost:9080/vivisimo/cgi-bin/admin`

To access the administrative interface from another computer, replace 'localhost' with the DNS name or IP address of this server.

You can start the web server manually by executing the **velocity-webservice start** command. Once enabled, the web server will be started automatically by the **velocity** command.

3. If the webservice has been enabled and is running, you will see a message like the following:

The webservice is enabled and running.

The url for the Velocity administrative interface is:
`http://localhost:9080/vivisimo/cgi-bin/admin`

To access the administrative interface from another computer, replace 'localhost' with the DNS name or IP address of this server.

- **install**: (Microsoft Windows only) Installs Watson Explorer Engine's embedded web server as a Window service.
- **restart**: Restarts the embedded web server by sending it a `SIGHUP` signal. If the web server process (`httpd`) is not running, it is started. This command automatically checks the configuration files via `configtest` before initiating the restart to make sure that the embedded web server will be able to restart with those configuration files.

- **start**: Starts the embedded web server that is included with Watson Explorer Engine.
- **stop**: Stops the embedded web server.
- **uninstall**: (Microsoft Windows only) Uninstalls the Watson Explorer Engine embedded web server as a Window service.

When executed with no arguments, the **velocity-webserver** command displays the same output as if it had been called with its **--help** option.

Messages

The **velocity-webserver** command displays the message "The webserver '*command*' is complete" when it is successful. Depending on the command-line options that you specify, you may be prompted for other information or see additional messages, such as the following:

- "The command must be one of start, stop, restart, graceful, configtest": You supplied an unknown keyword as an argument to the **velocity-webserver** command.
- "Press Return to continue": (Microsoft Windows only) Shown after the **velocity-webserver** completes the webserver startup process but before it exits, in order to enable you to see any messages that may have been generated in the command window during the startup process.

See the **Watson Explorer Engine Message Reference Manual** for detailed information about any other messages that you may receive from the **velocity-webserver** command.

Exit Status

Exits with status **0** when the command is successful.

Note: On Linux systems, exit codes above 127 are returned by the shell. Please consult the documentation for the shell that you are using for more information.

Integrating Watson Explorer Into the System Startup and Shutdown Process

Like other software that should always be running on a given computer system, you may want to integrate Watson Explorer and related modules into your system's startup and shutdown processes so that mandatory services are always running, and that those services are started in the right order during the startup process to satisfy dependencies, and are terminated in the right order during shutdown or reboot sequences to ensure data integrity.

Integrating Watson Explorer Engine Into the Startup/Shutdown Process

As explained in the next two sections, Watson Explorer Engine can easily be integrated into your system's startup and shutdown processes by integrating scripts that make calls to its command-line administration tools into those processes on your system(s).

Watson Explorer Engine Startup/Shutdown Integration on Linux Systems: The startup and shutdown commands that were discussed in the previous section can easily be integrated into the startup mechanism used by Linux systems by means of an init script, which is a file located in the `/etc/init.d` directory on those

systems, and which is typically invoked through symbolic links from various run-level-specific directories on those systems. Though different Linux systems provide system-specific mechanisms for handling init scripts, most of them provide compatibility with the traditional System V Init mechanism. Therefore, Watson Explorer Engine includes a sample init script, `velocity.sysvinit`, located in the `examples/init.d` directory of any Watson Explorer Engine installation. This script contains numerous comments that will help you customize it for your system's startup/shutdown mechanism. Without most of those comments, this script looks like the following:

```
#!/bin/sh
#
# You will have to set VIVISIMO_INSTALL_DIR to the location where
# Watson Explorer is installed on your system.
#
VIVISIMO_INSTALL_DIR=/opt/ibm/IDE/Engine
VNAME="Watson Explorer"

PATH=${VIVISIMO_INSTALL_DIR}/bin:/sbin:/bin:/usr/sbin:/usr/bin

check() {
    if [ ! -f ${VIVISIMO_INSTALL_DIR}/bin/velocity-startup ] || \
        [ ! -f ${VIVISIMO_INSTALL_DIR}/bin/velocity-shutdown ]; then
        echo " Could not find ${VNAME} admin commands in $VIVISIMO_INSTALL_DIR/bin - exiting."
        exit 1
    fi
}

case "$1" in
    start)
        check
        ${VIVISIMO_INSTALL_DIR}/bin/velocity-startup --yes \
            --start-scheduler --start-query-service | \
            grep -v '^$' | sed -e 's;^; ;g'
        ;;
    stop)
        check
        ${VIVISIMO_INSTALL_DIR}/bin/velocity-shutdown --yes | \
            grep -v '^$' | sed -e 's;^; ;g'
        ;;
    *)
        echo " Usage: $0 {start|stop}"
        exit 1
esac

exit 0
```

The only change that you may need to make to this sample script is to modify the value of the `VIVISIMO_INSTALL_DIR` variable, which specifies the location of the Watson Explorer Engine installation directory on your system.

Note: By default, this script executes the **velocity-startup** command in the way that most Watson Explorer Engine customers will want to use it, automatically starting the Watson Explorer Engine scheduler with a **boot** message, and automatically starting the Watson Explorer Engine query service. If you are not using the Watson Explorer Engine scheduler and/or do not want to start the Watson Explorer Engine query service automatically at boot time, you should remove the associated option(s) from the **start** case in the sample init script.

To integrate this script into your Linux system's startup and shutdown sequences, copy this script to the `/etc/init.d` directory on your system and create symbolic links to this script from the directories that are associated with the run levels in which you want to start Watson Explorer Engine. These are typically the `/etc/rc2.d`, `/etc/rc3.d`, or `/etc/rc5.d` directories, depending on your system's default run level. You should create both startup links (beginning with 'S') and shutdown links (beginning with 'K'). Suggested link names are `S90de_engine` and

K95de_engine. Before starting Watson Explorer Engine, the network, local and remote filesystems that are relevant to Watson Explorer Engine, and the web server with which Watson Explorer Engine is integrated must all be available. Similarly, Watson Explorer Engine should be shut down before any of these other services are shut down.

Note: Attempting to access a Watson Explorer Engine installation that has been shut down displays an error message stating that Watson Explorer Engine could not process your request because it is in maintenance mode. To resolve this problem, execute the **velocity-startup** command as the root user on a Linux system, or execute the **velocity-startup.exe** command on a Microsoft Windows system as a user with Administrator privileges.

Startup/Shutdown Integration on Microsoft Windows Systems: Though Microsoft Windows systems provide multiple mechanisms for integrating commands into various Startup folders, it is important to differentiate between commands that run at user login (personal startup commands), which can be placed in the Startup folder for that user (or for all users), and commands that run when a Windows system itself starts, regardless of whether anyone is logged in on the system. Commands such as Watson Explorer Engine's startup and shutdown commands should not be integrated into per-user startup mechanisms on production Watson Explorer Engine servers, because they would not start until a user logged in on those systems.

There are several mechanisms for integrating commands into the Windows system startup process. The most commonly-used of these are the following:

- **Group Policy Editor**, which is a Microsoft Management Console applet that enables you to add startup and shutdown scripts and associated parameters to the group policies for system startup and shutdown. This applet is a default part of Windows Server 2003, 2008, and 2012 installations. You should run this utility as a user with Administrator privileges to guarantee that you are able to modify basic system settings.

You can start this utility by selecting the Start menu's **Run** command, entering the command **gpedit.msc** in the **Open:** text area, and clicking **OK**. Once the tool starts, expand the **Computer Configuration** item in the left pane, expand the **Windows Settings** item in that same pane, and select the **Scripts (Startup/Shutdown)** item to display the **Startup** and **Shutdown** sections in the right pane. Double-clicking either of these opens an associated Properties dialog in which you can click **Add** to define a new startup or shutdown item. You should browse for and add the command **velocity-startup.exe** with the **--yes** option as a startup script, and add the command **velocity-shutdown.exe** with the **--yes** option as a shutdown script.

- **Registry Editor**, which has been available as a default part of Windows installation since Windows 3.1. You can add registry keys that will execute commands at system startup, but there are no default registry entries for executing commands at system shutdown. While it is possible, manually editing the registry to integrate Watson Explorer Engine's **velocity-startup.exe** command is not recommended.

You can start this utility by selecting the Start menu's **Run** command, entering the command **regedit** in the **Open:** text area, and clicking **OK**. Once the tool starts, navigate to the **HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion** section and add a key to the **RunServices** entry. The new key can have an arbitrary name, and its value should be the complete command-line that you want to execute, with all arguments. After adding this key, save your registry and exit the Registry Editor.

Once commands have been integrated into the system startup process, Microsoft Windows systems provide a number of different mechanisms for listing and enabling/disabling the commands that run at system startup. These mechanisms can be useful to temporarily disable commands that have been scheduled for execution at system startup, but do not provide a mechanism for adding new commands to the startup sequence. Examples of these commands are the following:

- **System Configuration Utility** (MSConfig), which has been provided as a default part of Windows installations since Windows 95. You can start this utility by selecting the Start menu's **Run** command, entering the command **msconfig** in the **Open:** text area, and clicking **OK**.

Once the System Configuration Utility starts, you can see and manipulate the applications that run at system startup by selecting the **Startup** tab.

- **Windows Defender**, which is primarily an anti-spyware/malware tool but also provides provided as a default part of Windows Vista and later systems. On Windows server systems, you can start this utility by selecting the Start menu's **All Programs** command, and selecting **Windows Defender** from that menu.

Once Windows Defender starts, select the **Tools > Software Explorer** menu item to see and manipulate the applications that are currently running on your system, and those that are scheduled for execution at system startup.

A number of third-party utilities provide alternate mechanisms for integrating commands into the Windows system startup and shutdown process. These have not been tested with the Watson Explorer Engine **velocity-startup** and **velocity-shutdown** commands, and are therefore not supported by IBM.

Integrating Other Watson Explorer Modules Into the Startup/Shutdown Process

As explained in the next two sections, the Application Builder, Results Module, Knowledge Center, and Zookeeper modules include services or startup scripts that you can easily integrate into your system's startup and shutdown sequence.

Startup/Shutdown Integration on Linux Systems:

Important: ZooKeeper must be running before you can successfully start and use Application Builder or any BigIndex API application. Applications that rely on ZooKeeper for configuration data will not start correctly if ZooKeeper is not yet available. Similarly, you must not terminate ZooKeeper before you terminate Application Builder - doing so may cause data loss because the Application Builder (or an application that was created with Application Builder) may not be able to correctly save all of its configuration data.

Startup and shutdown mechanisms on Linux platforms use descendants of a standard system startup and initialization mechanism known as **SysVinit**, or provide compatibility with it. SysVinit works by integrating run-level independent startup scripts into the startup and shutdown mechanism for any Linux runlevel. This is done by creating symbolic links from the runlevel-specific startup directory to the runlevel-independent script. The names of those symbolic links indicate the order in which the associated shell scripts are executed, whether those services are being terminated or started when entering or exiting a runlevel.

The startup/shutdown scripts that are installed when installing Watson Explorer and related modules contain comments that are used by the **chkconfig** command to automatically create symbolic links to those services which satisfy the dependencies between services such as ZooKeeper and Application Builder, where

ZooKeeper must be started before Application Builder so that Application Builder can load its configuration data from ZooKeeper.

The default names of the startup scripts that are installed in the `/etc/init.d` directory on Linux systems for Watson Explorer are:

- **appbuilder-service-default** - Application Builder
- **resultsmodule-service-default** - Results Module
- **knowledgecenter-service-default** - Knowledge Center
- **zookeeper-service-default** - ZooKeeper

After installing any Watson Explorer or related module from this list, you can integrate the associated startup scripts in the right order by executing a single administrative command for each such script. To do so:

1. Change directory to the `/etc/init.d` directory on your Linux system
2. For each script that you want to integrate into the system startup/shutdown process, execute the appropriate command for your Linux distribution as the root user or via the **sudo** command:
 - RHEL or compatible distributions:
`chkconfig --add SCRIPT-NAME`
 - SUSE or compatible distributions:
`chkconfig -a SCRIPT-NAME`

After integrating Watson Explorer services into the startup and shutdown sequence of a system, you should test that the system is configured correctly and that these services start and are usable whenever the system restarts. The easiest way to do so is by restarting the system.

Startup/Shutdown Integration on Windows Systems:

Important: ZooKeeper must be running before you can successfully start and use Application Builder or any BigIndex API application. Applications that rely on ZooKeeper for configuration data will not start correctly if ZooKeeper is not yet available. Similarly, you must not terminate ZooKeeper before you terminate Application Builder - doing so may cause data loss because the Application Builder (or an application that was created with Application Builder) may not be able to correctly save all of its configuration data.

As part of installing Watson Explorer on a Microsoft Windows system, the Watson Explorer installation process also installs related services on your Windows system. These services enable you to automatically start Watson Explorer (and related modules, if necessary) in the correct order on that system as part of the Windows startup process, and to terminate those services in the correct order when shutting down or restarting that system.

The services that are installed on Windows system for use by Watson Explorer and related modules are initially configured to start manually, though they create the right dependencies. Before permanently enabling them, you should start them manually to ensure that services start in the right order, and with sufficient time between service starts to ensure that services are fully started by the time another application requests information from them. See “Startup/Shutdown Integration on Microsoft Windows Systems” on page 77 for a general discussion of services on Microsoft Windows systems.

Integrating Watson Explorer with Anti-Virus Software

Many systems today run background processes or system services, such as anti-virus, anti-spyware, and anti-malware packages, that regularly scan the filesystems on a computer system. These packages check new and updated files for known signatures of viruses and other problematic software. On Microsoft Windows systems, the executable files produced by IBM as part of Watson Explorer are all digitally signed to identify their origin and to verify their contents.

If you are running such scanning software, make sure that you exclude the directory where you installed Watson Explorer (c:\Program Files\IBM\IDE\ on Microsoft Windows systems, and /opt/ibm/IDE on Linux systems), as well as any other location(s) in which you are storing data (crawl, index, and configuration data), from the list of directories to be scanned. Index files created by Watson Explorer are large binary files, and cannot contain viruses or dangerous software. They can also change relatively frequently, depending on the index refresh mode that you select or the frequency with which you manually update your index files. Scanning these files will hurt the performance of your system without providing any potential benefit. Anti-Virus scanning can also permanently corrupt an index by attempting an automatic replacement of index code that happens to match the signature of a known virus.

If your system security policies require that all executable binaries and libraries must be scanned for infection, you can simply exclude the following sub-directories of your Watson Explorer installation directory, as well as any other directories in which you are storing data:

- Engine/cache
- Engine/data
- Engine/debug
- Engine/tmp
- ResultsModule/wlp/usr/servers/ResultsModule/apps/ResultsModule/WEB-INF/db
- AppBuilder/zookeeper-data

These locations contain indices, collections, configuration data, applications, and temporary files produced and used by Watson Explorer, but do not contain any binary applications or libraries.

In addition to the negative performance impact of scanning files that do not need to be searched for viruses or malware, not excluding your Watson Explorer directories from virus scanning can also prevent standard Watson Explorer operations, such as pushing staging indices to live indices and cleaning temporary files, from operating correctly. Most anti-virus software keeps files open as they are being scanned, which will prevent Watson Explorer from deleting those files if it needs to do so.

Integrating Watson Explorer Engine with System Cleanup Scripts

Watson Explorer Engine uses named pipes to communicate between the **collection-service** process and the associated search collection. These named pipes are normally located in the directory that holds the index and other files for a search collection, but are stored in your system's temporary directory (/tmp on Linux systems) if the filename exceeds 128 characters. Filenames for the named

pipe files that are created and used by Watson Explorer Engine are of the form `/tmp/Vivisimo Velocity - * or /tmp/viv_conv_pipe_*`.

Many systems automatically run applications such as `tmpwatch` or use custom scripts that clean up system temporary directories at scheduled intervals. Any such scripts that are running on systems where Watson Explorer Engine is installed must be modified to prevent the removal of named pipes located in system temporary directories. Watson Explorer Engine will automatically clean up the named pipe files that it creates after it is done using them.

To proceed to the next section, which discusses integrating Watson Explorer Engine with automated system update software, click “Integrating Watson Explorer Engine with System Updates.”

Integrating Watson Explorer Engine with System Updates

Many computers today run system services that contact an operating system vendor and automatically download and install system updates. These updates typically provide security-related software updates, bug fixes, and any package enhancements or new packages required by those updates.

Installing updated operating system binaries or system libraries may require rebooting your system in order to remove old versions and begin using the new ones. Some software update systems automatically reboot your system when it is necessary to do so after a software update. If this is the case, you must make sure that software updates are scheduled so that they cannot conflict with Watson Explorer Engine index updates or re-crawls of your search collections. Rebooting the system while an index update is in progress can substantially increase the time required for index creation or updates, especially for indices that are being created by crawling remote resources that may not always be available or to which connectivity is intermittent or slow.

To proceed to the next section, which discusses command-line tools for Watson Explorer Engine administration, click “Command-Line Watson Explorer Engine Administration Commands” on page 68.

Backing Up Watson Explorer Engine

You can always back up your Watson Explorer Engine installation by adding its installation directory to the list of standard directories that you back up. The default location of the Watson Explorer Engine installation directory is `/opt/ibm/IDE/Engine` on Linux systems, and is `C:\Program Files\ibm\IDE\Engine` on Microsoft Windows systems, though the installation directory can be changed when installing the software.

warning: You must not back up the Watson Explorer Engine installation directory while any index updates are taking place. This is primarily a scheduling issue - make sure that index updates are only scheduled to occur after backups have completed. If indices are being updated while they are being backed up, they may not be usable after being restored. Similarly, if your backup application locks files before backing them up, Watson Explorer Engine may terminate abnormally because it cannot write to the files and directories that it requires.

At a minimum, you will want to add the files `data/users.xml` and `data/repository.xml` to your daily backups. The `users.xml` contains the user accounts for your Watson Explorer Engine installation. The **repository.xml** file contains the configuration code for all of your search applications.

Note: If you are developing collaborative search applications and have enabled annotation backups as described in the *Express Tagging* section of the *Watson Explorer Engine User Manual*, you will also want to add the directory where these automatic backups are stored to your system backups. The backups for search collection annotations are stored in the directory `data/tag-backup`, relative to the directory where Watson Explorer Engine is installed on your system. The backup file for search collection annotations is updated each time any annotation in that collection is added or modified, so you may want to back up this directory more frequently than you perform standard system backups. This is easily done by creating a shell script or batch file that runs at scheduled intervals and updates remote copies of the files in this directory if they have been modified. For information about restoring annotations from a backup file, see *Restoring Annotations from a Backup File*.

Beyond these critical files, deciding what to back up depends on how you are using Watson Explorer Engine and the types of search applications that you are creating. For example, if your Watson Explorer Engine applications are primarily meta-search applications and therefore do not themselves crawl or index other sources of online information, you may not want to back up anything beyond the files mentioned earlier in this section.

However, if you are crawling data repositories to create your own search collections, all of the indices, log files, and other data from all sources that you are crawling will also be stored under your Watson Explorer Engine installation directory by default. The crawled and indexed data for a Watson Explorer Engine search collection will be stored in the directory `data/search-collections/XYZ/name`, where *name* is the name of the search collection, and *XYZ* is the first three bytes of an internal hash that was calculated from that name. Storing search collections in distinct directories helps prevent directory size and performance issues.

Important: Prior to Velocity 7.5, all search collection data was stored by default in the `data/collections` directory of an installation. Any search data for search collections that you created prior to Velocity 7.5 in an installation that you have upgraded to Velocity 7.5 will still be located in the `data/collections` directory of the installation. This directory should therefore also be backed up if you want to back up those collections, taking into account the same considerations discussed in this section for backing up any search collection.

Note: Search collection data location is configurable on a per-collection basis in the **Directories** section of the **Configuration > Meta** tab for that search collection in the Watson Explorer Engine administration tool. If you manually configure the location where the data for any search collections is stored, you will also have to back up those directories if you are backing up your search collection data.

The files that make up a search collection can be very large, depending on the amount and type of data that you are crawling. You will need to make the traditional system administrator's decision regarding the trade-offs between increased backup time and backup storage requirements compared to the time it would take for Watson Explorer Engine to re-crawl the data and recreate the indexes (or update them if you have restored them from backups).

Note: If you are using Watson Explorer Engine to crawl and index data, the index and log files that it creates and uses are large, binary files. Therefore, it is not really possible to make incremental backups of these files with any granularity less than the frequency with which they are updated by new crawls or index updates. Most backup software uses file creation and modification times to determine which files

have changed and thus need to be backed up. The timestamps of Watson Explorer Engine's log and index files will be updated each time that index data is actually updated in response to a recrawl or refresh request. (Refresh requests will only actually update an index if data has changed on the resource that is being indexed.) Some backup software supports saving only the changes between two versions of a file, but binary files such as these will usually appear to be 100% different than previous versions of these files.

In general, if re-indexing your search collections does not take very long (or takes less time than it takes to restore selected files from backups), you may not want to back up anything other than the files mentioned earlier in this section.

If re-indexing your search collections takes a significant amount of time but is done infrequently, you will only need to back up your Watson Explorer Engine installation directory after each time that you re-index your search collections. This meshes nicely with standard incremental backup mechanisms, which only back up files that have changed since a given date or since backups were last run.

If re-indexing your search collections takes a significant amount of time but is done daily, you will need to determine, on average, how long it takes to re-crawl your information sources and update an existing index. If daily updates take a relatively small amount of time, you may want to back up your Watson Explorer Engine installation directory weekly and, whenever necessary, update them after they have been restored. If the time that it takes to restore indices from backups and update them approaches the amount of time it takes to create them, you may not want to back them up at all.

IBM strongly recommends installing your Watson Explorer Engine software on RAID storage so that the failure of a single disk will not cause your Watson Explorer Engine applications to fail. If your applications require high availability, you can take advantage of Watson Explorer Engine features such as *Distributed Indexing* to help protect against system failures. You may also want to consider co-location to protect against local or regional power or network failures.

Tip: For general information about hardware/software requirements and minimum system configurations, see "Requirements" on page 1.

Notices and trademarks

This information was developed for products and services offered in the U.S.A. This material may be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

Notices

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 grant 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 character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 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 may not apply to you.

This information 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 serve 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
J46A/G4
555 Bailey Avenue
San Jose, CA 95141-1003 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-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their 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.

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 similarity 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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice 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 softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at www.ibm.com/legal/copytrade.shtml.

The following terms are trademarks or registered trademarks of other companies:

Adobe is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

The United States Postal Service owns the following trademarks: CASS, CASS Certified, DPV, LACS^{Link}, ZIP, ZIP + 4, ZIP Code, Post Office, Postal Service, USPS and United States Postal Service. IBM Corporation is a non-exclusive DPV and LACS^{Link} licensee of the United States Postal Service.

Other company, product or service names may be trademarks or service marks of others.

Providing comments on the documentation

You can provide comments to IBM about this information or other documentation.

About this task

Your feedback helps IBM to provide quality information. You can use any of the following methods to provide comments:

Procedure

- Send your comments by using the online readers' comment form at www.ibm.com/software/awdtools/rcf/.
- Send your comments by e-mail to comments@us.ibm.com. Include the name of the product, the version number of the product, and the name and part number of the information (if applicable). If you are commenting on specific text, include the location of the text (for example, a title, a table number, or a page number).



Printed in USA

GC19-4189-05

