



Lotus Expeditor 6.1 Education

IBM® Lotus® Expeditor 6.1 Server Clustering Overview

Lotus software



@business on demand software

© 2006 IBM Corporation

Welcome to this overview of the IBM Lotus Expeditor 6.1 Server Clustering. You will see an overview of clustering including design goals and requirements. You will also walk through the high level steps to creating a cluster.

Clustering overview

- Start with single box remote database install
- Configure for Active Directory or custom LDAP user registry
- Configure Node and Cell for managed environment
- Cluster first node
- Add additional cluster nodes

Clustering starts with a single box install. This allows for installs to start small and grow. It also allows for the initial functionality to be proven before committing resources to a clustered deployment. The initial install must start with a remote database install. There is no migration or configuration support to change from a local database install to a remote database install.

After the initial remote database install is installed and configured, the node goes through a set of configuration steps to become clustered. Clustering requires an LDAP user registry. You can use the configuration wizard to upgrade to the Active Directory 2003 user registry, or you can use command-line tools to upgrade to a custom LDAP.

Clustering application servers requires a managed environment, but the initial install is only supported into a non-managed environment. The configuration wizard provides steps to configure the cell for the same security used by the primary node. Security can already be configured for the cell but it must match the security settings used by the primary node. The configuration wizard will also add the node to the managed environment. It is required that the configuration wizard add the node to the managed environment or settings will be lost.

Once the node is added to the managed environment, the configuration wizard will create the initial clusters. It will also generate a cluster response file that can be used to add additional nodes to the cluster.

Adding additional cluster nodes is done by an option in the Launchpad. That option launches the install with special parameters indicating that this is an additional node. This install uses the cluster response file from the primary node to simplify gathering information about the primary node.

Clustering design goals

- Start small and grow
- Establish that each step is functional before continuing
- Response file from first node allows simplified install on additional nodes

The clustering design allows the user to start with a single box install and grow that install into a cluster later. The initial install is fully functional and allows for validation before continuing. Once the first node is clustered, there is a cluster response file generated that simplifies installs of additional nodes.

Clustering requirements

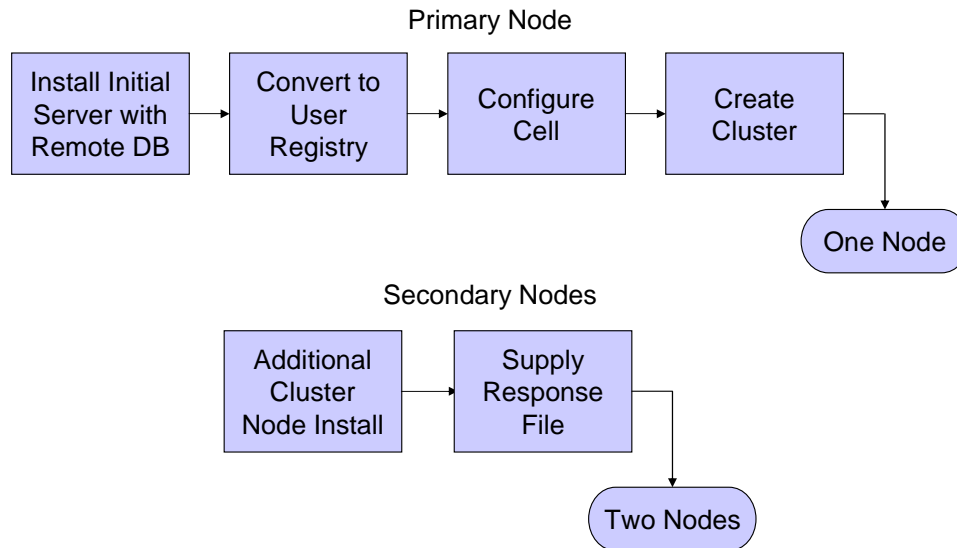
- Primary Node
 - ▶ Security cannot be enabled in the profile
 - ▶ Cannot be a managed node (part of a cell)
 - ▶ Must be added to the cell by the Configuration Wizard
- Linux®
 - ▶ Requires a shared group containing both root and the DB2® instance user

Since clustering starts small, the primary node must start with the same requirements as the single box install. It cannot have security enabled and it cannot be part of a cell.

The process of adding the single box install to a managed environment requires several steps. These steps are contained in the **Configure Node for Cell Environment** task in the configuration wizard. This task will add the node to the cell and it will ensure that all the necessary applications and settings are transferred into the cell as well. If you add the node to the cell manually, the node will no longer work.

On Linux, there is an added special requirement. Some of the application servers run as the DB2 instance user while others run as the root user. To maintain shared access to certain files, the install needs a shared group that contains the root user and the DB2 instance user.

Clustering flow



You can see here the two steps involved in growing into a cluster.

The primary node path starts with the initial install with a remote database. After the install is complete, the server is upgraded to use a LDAP User Registry instead of the local database user registry used initially. For Active Directory, this can be done using the configuration wizard. For other LDAPs, there are command line tools available.

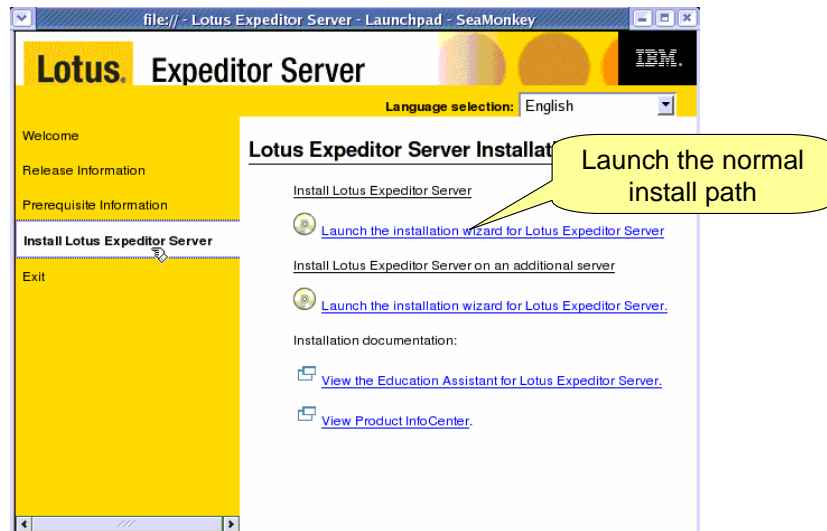
After that step, the configuration wizard is used to configure the node for the cell environment. This includes confirming the security of the node and the cell match. If the node was upgraded to Active Directory using the configuration wizard, the configuration wizard will configure security on the cell for you. If you used a custom LDAP, then you must manually configure the cell security to match the local node security. Once the security is validated, this step adds the node to the cell.

Finally there is a configuration wizard task to create the initial cluster. During this process a cluster response file is created. That file must be transferred to each new node as input to installing the secondary node.

For the secondary nodes, there is an option on the launchpad to install additional servers. This install will prompt for the response file and additional information about the node installed on the secondary server. The secondary server must have WebSphere and DB2 already installed. IBM HTTP Server is not required on secondary nodes. Unlike the primary node, the secondary node can already be a member of the same managed cell to which the primary node belongs.

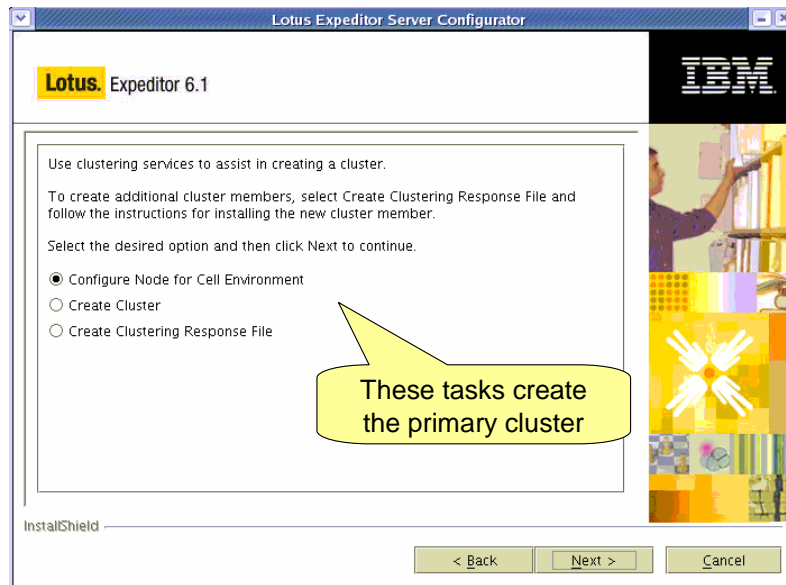
It will then install the files and configure this new node as an additional cluster member. This includes adding this new node to the cell if the node is an unmanaged node.

Clustering – Primary node launch



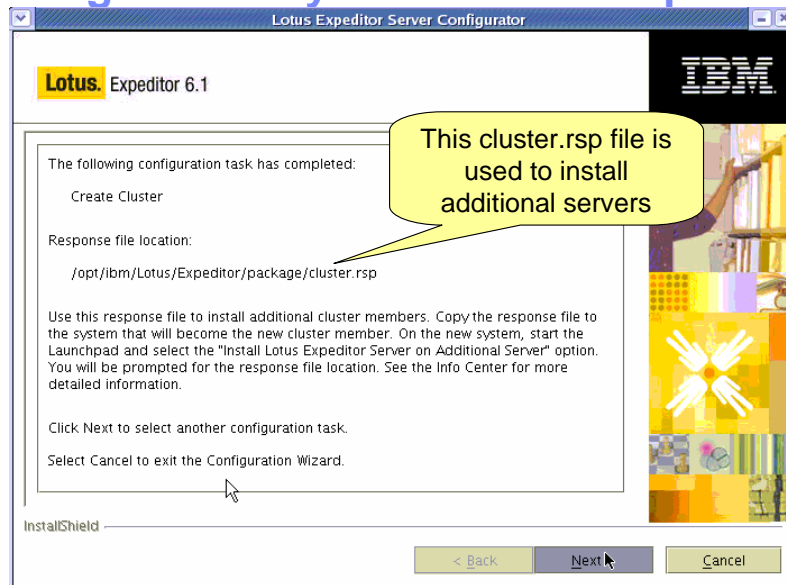
Starting to create a cluster is the same as the normal install path. This slide shows the launchpad and the two options for installing the server. The first option is the most common and used for normal installs. The second option is only used when a cluster already exists.

Clustering – Primary node configure



After the install is completed, the configuration wizard provides simple options to step through the clustering process. On this panel you would first configure the node for the cell environment, then you would create the initial cluster.

Clustering – Primary node cluster.rsp



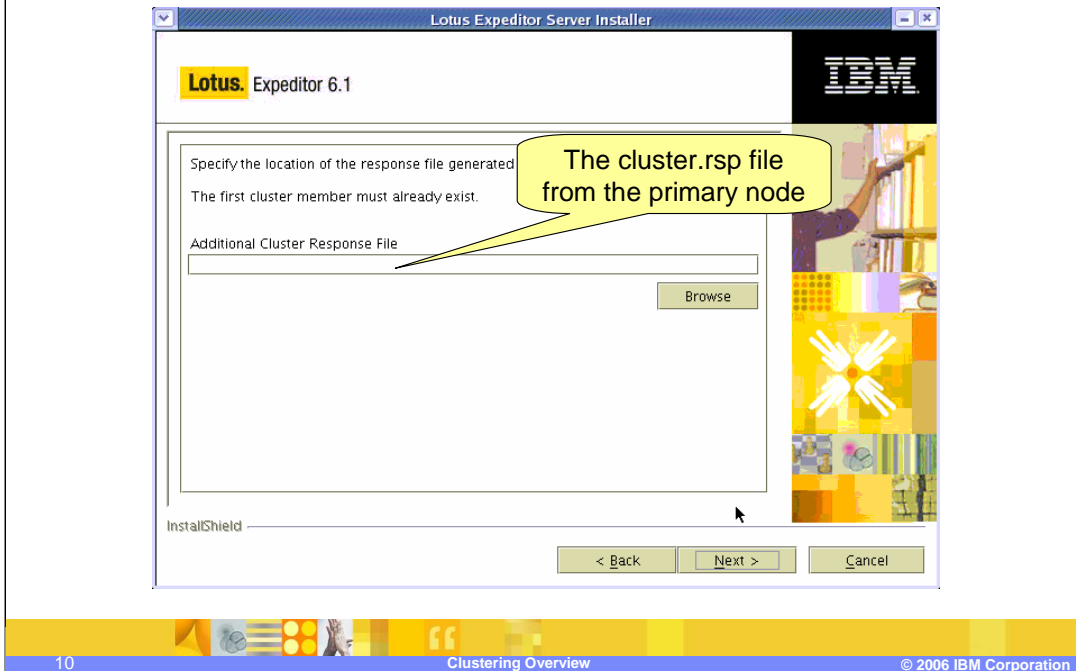
Once the cluster is created, you should see a summary panel like this. Note the location of the cluster.rsp file. This is the file used to transfer settings from the primary node to the secondary nodes. Note that this file does not contain any passwords. The secondary node installs will prompt for this response file.

Clustering – Secondary node launch



Installing secondary nodes is done by selecting the second option in the launchpad. This option launches the same installer but with slightly different parameters so that it knows to configure this node as an additional server.

Clustering – Secondary node cluster.rsp



During the additional server install, it will prompt for the location of the cluster.rsp file generated from the primary node. Then it will prompt for some information specific to the new node. Once this install completes, the cluster will have one additional member.

Clustering summary

- Start small
- Reconfigure to cluster
- Add additional nodes

To summarize the clustering process:

It starts small with a single box using a remote database.

The single box is reconfigured to create the initial cluster.

Additional nodes can be added using an install with a response file.

Trademarks, copyrights, and disclaimers

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

DB2 IBM Lotus

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

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements or changes in the products or programs described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

Information is provided "AS IS" without warranty of any kind. THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (for example, IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. 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 in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products.

IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

© Copyright International Business Machines Corporation 2006. All rights reserved.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.

This concludes this presentation.