

HP OpenView Customer Views for Network Node Manager

Configuring Customer-specific Network Management

HP-UX, Solaris, and Windows



i n v e n t

Manufacturing Part Number : J4518-90012

May 2003

© Copyright 1999-2003 Hewlett-Packard Development Company, L.P.

Legal Notices

Warranty.

Hewlett-Packard makes no warranty of any kind with regard to this manual, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be held liable for errors contained herein or direct, indirect, special, incidental or consequential damages in connection with the furnishing, performance, or use of this material.

A copy of the specific warranty terms applicable to your Hewlett-Packard product can be obtained from your local Sales and Service Office.

Restricted Rights Legend.

Use, duplication or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause in DFARS 252.227-7013.

Hewlett-Packard Company
United States of America

Rights for non-DOD U.S. Government Departments and Agencies are as set forth in FAR 52.227-19(c)(1,2).

Copyright Notices.

No part of this document may be copied, reproduced, or translated to another language without the prior written consent of Hewlett-Packard Company. The information contained in this material is subject to change without notice.

©Copyright 1999-2003 Hewlett-Packard Development Company, L.P., all rights reserved.

Contains software from AirMedia, Inc.

© Copyright 1996 AirMedia, Inc.

Trademark Notices.

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

Java™ is a U.S. trademark of Sun Microsystems, Inc.

Microsoft® is a U.S. registered trademark of Microsoft Corporation.

Windows® and MS Windows® are U.S. registered trademarks of Microsoft Corporation.

Netscape™ and Netscape Navigator™ are U.S. trademarks of Netscape Communications Corporation.

Oracle® is a registered U.S. trademark of Oracle Corporation, Redwood City, California.

Oracle7™ is a trademark of Oracle Corporation, Redwood City, California.

OSF/Motif® and Open Software Foundation® are trademarks of Open Software Foundation in the U.S. and other countries.

UNIX® is a registered trademark of The Open Group.

All other product names are the property of their respective trademark or service mark holders and are hereby acknowledged.

1. Integrating Customer Data

Overview of Customer-specific Network Management	10
Theory of Operation	10
The ovcustomer Program	11
Organization Groups	12
Configuring Customer-specific Network Management	14
Integrating Your Customer Data	14
Method 1: Execute a Batch File	15
Method 2: Manually Issue ovcustomer Commands	18
Creating Key Resources	19
Identifying Customer Equipment	20
Creating Locations	21
Using an External Customer Database	22
For More Information	24

2. Using Customer Views on Management Consoles and in Distributed Environments

Using Customer Views on a Management Console	27
Licensing	27
Installation	28
Using Customer Views in a Distributed Environment	28
Managing Customers from the Management Station	28
Managing Customers from the Management Station and Collection Stations	29
Best Practices	31

Support

Please visit the HP OpenView web site at:

<http://openview.hp.com/>

There you will find contact information and details about the products, services, and support that HP OpenView offers.

The support area of the HP OpenView web site includes:

- Downloadable documentation
- Troubleshooting information
- Patches and updates
- Problem reporting
- Training information
- Support program information

1 Integrating Customer Data

This chapter describes the Customer-specific Network Management feature of HP OpenView CustomerViews for NNM and how to integrate customer information into NNM's database.

Overview of Customer-specific Network Management

HP OpenView Network Node Manager provides many mechanisms for managing network resources. Customer Views for NNM helps you manage those network resources in the context of the specific customers who depend on them.

Customer Views for NNM enables you to monitor and troubleshoot customer links and Internet access links, associate customers to the network resources they utilize, group customers based on attributes, and determine which customers are impacted by problems on the network.

Understanding and responding to network events based on which customers, providers, or peers are affected enables you to improve service, and increase customer retention and satisfaction.

You can perform the following tasks using Customer-specific Network Management:

- Create, configure, and modify customer relationships.
- Collect data and produce reports on specific customer links.
- Determine which customer or provider is impacted when a network resource goes down.
- Associate customers with the network resources they use.
- Monitor performance of links for specific customers.

Theory of Operation

Customer Views for NNM uses the `ovcustomer` program to help you manage your customers in relation to the network resources they use. In addition, Customer Views allows you to group your customers into organization groups, which further enhances your management capabilities.

The ovcustomer Program

NNM has a powerful engine for automatically discovering the elements in your network and determining how they are interconnected. However, NNM cannot automatically determine who your customers are and which resources they are using. To provide this information to Customer Views, you use the `ovcustomer` program.

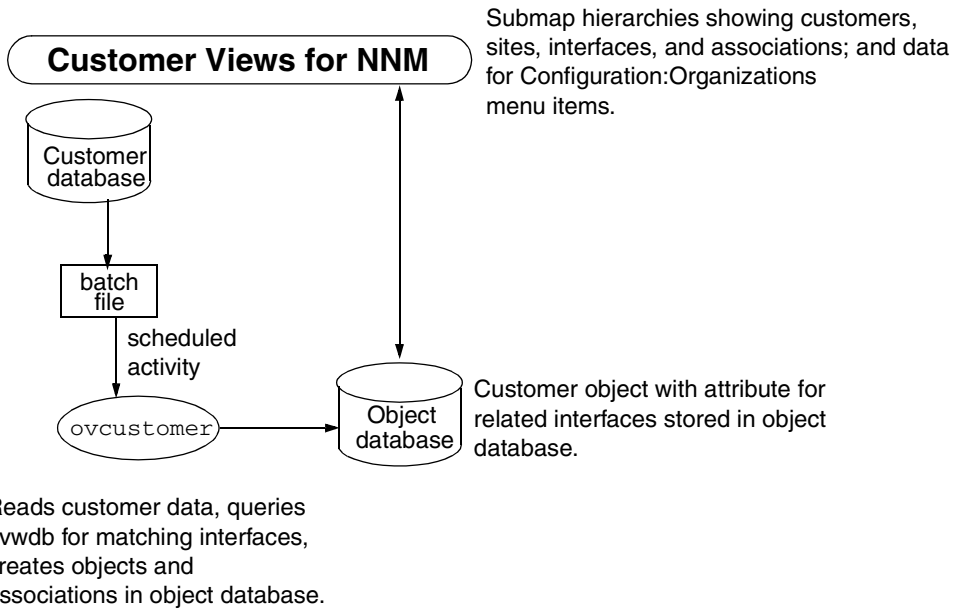
The `ovcustomer` program, at its most basic level, allows you to create and delete customers, and associate and unassociate network resources with them. This information is stored in NNM's database, `ovwdb`. `ovcustomer` also allows you to view and alter this information; however, Customer Views' interface and associated dialog boxes probably provide a more suitable mechanism for viewing this data. Once customers have been created and associated with network resources, the submap shows these customers and allows the user to navigate among their associations.

Typically, `ovcustomer` will be used in a batch mode to load a large set of information from an external database. Its interactive mode is useful when first exploring the functionality of the product, and for interactively modifying customer information once it has been added to NNM's database.

Once the customers have been loaded into the NNM database, they will appear in the Customer submap, accessible from NNM's root map. Each customer appears as an explodable symbol that shows the network resources associated with that customer.

Figure 1-1 shows how information flows from your customer database to HP OpenView's object database, and how that information is used.

Figure 1-1 Information Flow in Customer Views



Although NNM's database can manage many properties of organizations, its primary purpose is for the efficient operation of NNM and for managing the data it presents. Therefore you will get the best results if you use NNM's database to manage a minimal set of customer information, and store the majority of customer information in an external database. Then those two databases can work in tandem. Details for interaction with an external database are on page 22.

Organization Groups

Customer Views for NNM uses the concept of organization groups to help you manage customers, providers, and peers. Typically the number of customers to whom you are providing service will be large, while the number of providers will be small. Using organization groups is especially useful when you have a very large number of customers, as you can provide some logical segmentation of the customers, which reduces the cognitive load on your operators.

Organization groups may also help increase system performance. For example, if you have more than a few hundred customers, displaying the Customer submap may be sluggish and an operator will find such a submap unusable simply due to the number of symbols on it. However, if you divide these customers into some logical groups where the groups have less than 100 customers, the submap performance will be acceptable and the operator can more easily find useful information and focus on the most important customers. A typical grouping might be by class of service.

Your operators may want to concentrate their attention on a few select customers, especially if the number of customers is quite large. In this case, you may find it useful, to leave some customers ungrouped (that is, not in any group) and to set an option to have these ungrouped customers not shown on NNM's submaps. These ungrouped customers can still be accessed via other dialog boxes in Customer Views.

Configuring Customer-specific Network Management

The import/query utility, `ovcustomer`, can be used in several ways:

- From the command line, with parameters.
- Interactively, from standard input (`stdin`).
- To execute a batch file containing customer information.

When first becoming acquainted with this functionality you may find it helpful to use `ovcustomer` interactively to issue individual operations (Method 2 below), or you may find it helpful to generate a small batch file by hand (Method 1 below).

Eventually you will want to produce a utility to extract data from your current customer database and automatically generate a batch file for processing by `ovcustomer`. This utility should be run periodically on your customer database to keep it in synch with the object database. If your database is an SQL database you may want to issue SQL statements, which can directly generate the appropriate statements for processing by `ovcustomer`.

The type of organization (customer) information that `ovcustomer` can integrate into the object database includes:

- Organization name
- Organization ID # (optional)
- Organization group (optional)
- Organization type (provider, customer)
- Associations between organizations and nodes
- Associations between organizations and interfaces

Integrating Your Customer Data

There are two methods for integrating your customer data:

- Method 1: Enter your customer information into a file, then run `ovcustomer` on that file.

- Method 2: Manually issue `ovcustomer` commands

Below is an example of a simple customer management database. Data from this sample database is used in the examples that follow this section.

SQL databases are constructed with data in tables, with each table usually having multiple columns. For example, one table in your database may contain the customer name and an ID number that you have assigned to that customer. Another table may contain the customer name, the name of the host or server that the customer uses, and the IP addresses of the interfaces on which the customer connects to your network.

To simplify these examples, assume that all your customer data is in one table, which might look something like the following:

Table 1-1 Customer Database Information

Name	External Key (ID)	Host/server	Interface1	Contact	email
UniComm	A-476	ace.xisp.com	12.3.156.78	D. Dennis	dd@ace.com
State Gas	C-158	serv2.xisp.com	13.4.178.9	W. Wolfe	ww@stgas.com
Metro Power	M-946	admin.com	23.7.123.32	M. Moore	mm@metrop.com

1 The Interface name in your table must be mappable to an IP address. You do not have to exit NNM to perform these commands.

Method 1: Execute a Batch File

The first method of integrating your customer data into NNM's database is to create a batch file that you direct `ovcustomer` to execute. The steps you will perform are:

1. Enter your customer database information into a batch file.
2. Execute the batch file with `ovcustomer`.
3. Verify that the information was added to NNM's database.

The following examples use the sample customer database information in Table 1-1. We will not include the contact name or email address from the sample database, as these fields are not currently supported by Customer Views. These fields can be accessed through your external customer database, described on page 22.

Step 1: Create customer information batch file

You need to create a batch file that contains the names of your customers, their ID numbers, nodes they use, and interfaces assigned to them. You can name this file anything you want—for our example, let's call this file `customer_info`. Using the tables above, the file would look like this:

```
# Create our customer list
create_organization customer "Central Supply" C-156
create_organization customer "Metro Power" M-946
create_organization provider "UniComm" A-476

# Add these hosts and associate them with our customers
add_associations_to_org "Central Supply" serv2.xisp.com
add_associations_to_org "Metro Power" admin.com
add_associations_to_org UniComm ace.xisp.com

# Add these interfaces and associate them with our customers
add_associations_to_org "Central Supply" 13.4.178.9
add_associations_to_org "Metro Power" 23.1.123.32
add_associations_to_org UniComm 12.3.156.78
```

Optionally you can add the following statements, which group your customers into organization groups, even though this information is not part of your customer database.

```
# Let's create some groups and assign our customers to them
set_organization_group "Central Supply" Utilities
set_organization_group "Metro Power" Utilities
```

Note that quotes are required when there is a space in a field, such as in the customer names above.

Alternative Method for Creating the Batch File

Instead of creating the batch file `customer_info` manually, you can use SQL statements to extract the information from your customer database. If the tables are formatted similar to the example tables shown above, you can use this method.

1. To merge the information from Table 1-1 above into the file, enter the following:

```
select "create_organization", "customer", 'Name', 'External Key' from  
;table_name
```

where *table_name* is the name of the table in your SQL database that contains this information.

2. Add the node and interface information from Table 1-1 above to the file. At the command line, enter the following:

```
select "add_associations_to_org", 'Name', 'node' from table_name;  
select "add_associations_to _org", 'Name', 'interface' from table_name;
```

where *table_name* is the name of the table in your SQL database that contains the node and interface information.

Step 2: Add customer information to the database

Once the batch file is created, the next step is to execute the file, which merges the information from the `customer_info` file into NNM's database. At the command line, type the following:

```
ovcustomer < customer_info
```

This command directs `ovcustomer` to read the file `customer_info`, and adds fields for the customer name, external key, node, and interfaces to NNM's database. If you added the `set_organization_group` statements to the batch file, the organization groups will be created and the customers added to them.

Step 3: Verify that the information was added to the database

You can verify that the information was added to NNM's database in several ways:

- Use one of the following `ovcustomer` commands:

```
ovcustomer print_organization
```

This command lists organization name, organization type, organization external key, and organization group.

ovcustomer export_all_organizations

This command lists the entire set of organizations and their associations with nodes and interfaces.

- Start NNM and view the information in the Customer submap hierarchy.
- If NNM is already running, the information should now appear in the Customer submap hierarchy.

Method 2: Manually Issue ovcustomer Commands

You can use the `ovcustomer` command to manually add information about individual customers to NNM's database. Following are some examples, using both the command line method and standard input.

Example 1

Add UniComm, which is a provider, to the object database (using the command line interface):

```
ovcustomer create_organization provider UniComm
```

Example 2

Associate the interface of IP address 12.3.156.78 with UniComm (using the command line interface):

```
ovcustomer add_associations_to_org UniComm 12.3.156.78
```

Example 3

Create the customer Metro Power and associate a node with Metro Power (using interactive mode):

```
ovcustomer
add_associations_to_org "Metro Power" serv2.xisp.com
quit
```

Example 4

Assign the customers Metro Power and State Gas to the organization group called Utility (using interactive mode):

```
ovcustomer
set_organization_group "Metro Power" Utility
set_organization_group "State Gas" Utility
```

Note that you do not have to create the organization group first; it is automatically created at the same time you assign a customer to it. The result of this command is that the customers Metro Power and State Gas now appear in the submap hierarchy in a submap called Utility.

All other customers not assigned to a group now appear in a submap called Other.

Example 5

This is an alternative method for creating an organization group and assigning customers to it. You set a default organization group that all subsequently-created customers will be assigned to. This example assumes that Metro Power and State Gas have not yet been created.

```
ovcustomer  
set_default_org_group Utility  
create_organization customer "Metro Power"  
create_organization customer "State Gas"
```

The result of this command is the same as in Example 4. However, the default group that you set here remains active only for this invocation of `ovcustomer`.

Creating Key Resources

Key resources are those network devices, such as servers or interfaces, and customers and providers that are critical to the management of organizations. You can identify which of these sites, equipment, etc. you consider key resources. Then you can manage those key resources through their own submap.

There are two ways to identify key resources:

- Through the `ovcustomer` command.
- Through the object properties dialog box.

To use the `ovcustomer` command, you use the `set_field` option to specify that an object (server, interface, etc.) is a key device. Following is an examples using the devices that were used earlier.

```
ovcustomer set_field admin.com isKeyDevice true
```

To identify key resources through the object properties dialog box, follow these steps.

Configuring Customer-specific Network Management

1. Select the device on the submap that you want to be a key resource.
2. Click the right mouse button to display the symbol popup menu.
3. Select `Object Properties` from the popup menu. The `Object Properties` dialog box appears.
4. Select `General Attributes`, then press `[Edit Attributes]`.
5. To specify that this device is a key resource, select the field `isKeyDevice`. The value changes from “False” to “True”. Then press `[OK]`.
6. Press `[OK]` to exit this dialog box.

Identifying Customer Equipment

You can identify certain network resources as customer premises equipment (CPE). These resources then appear as part of that customer’s submap, allowing you to easily associate them with the customer.

There are two ways to identify customer premises equipment:

- Through the `ovcustomer` command.
- Through the object properties dialog box.

Make sure that you first, associate the network resource with the customer.

```
ovcustomer add_associations_to_org UniComm admin.com
```

Then use the `set_field` option to specify that an object (server, interface, etc.) is set to `isCPE`.

```
ovcustomer set_field admin.com isCPE true
```

To identify customer premises equipment through the object properties dialog box, follow these steps.

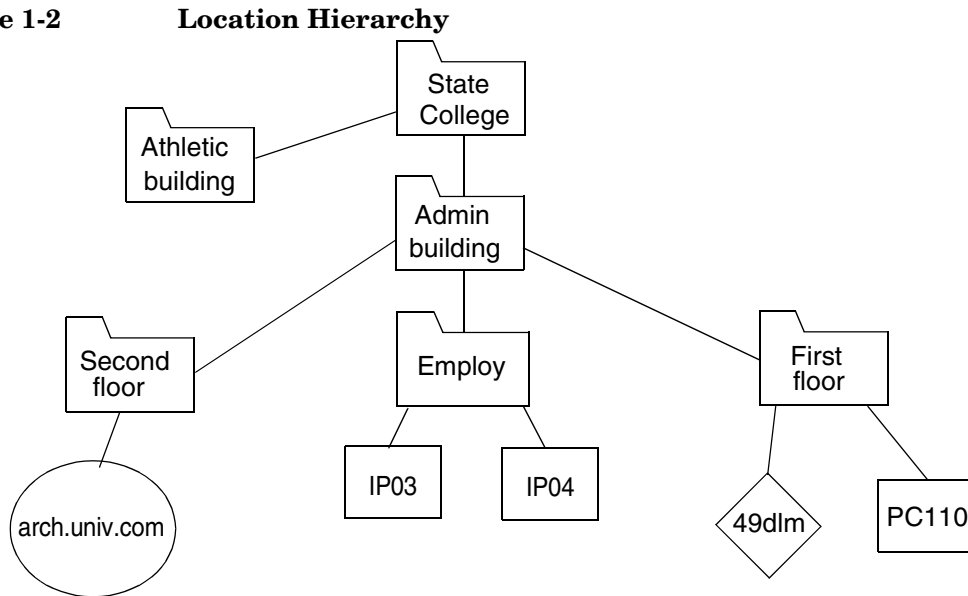
1. Select the device on the submap that you want to be a key resource.
2. Click the right mouse button to display the symbol popup menu.
3. Select `Object Properties` from the popup menu. The `Object Properties` dialog box appears.
4. Select `General Attributes`, then press `[Edit Attributes]`.

5. To specify that this device is customer premises equipment, select the field `isCPE`. The value changes from “False” to “True”. Then press [OK].
6. Press [OK] to exit this dialog box.

Creating Locations

Setting up locations is a convenient way to organize networks, nodes, and devices, such as interfaces, switches, and routers. This helps you to monitor key service provider locations. Figure 1-2 on page 21 shows an example of a simple hierarchy. Directions for setting up this hierarchy using `ovcustomer` follow the figure.

Figure 1-2



Create an organization hierarchy similar to that shown in Figure 1-2.

```
ovcustomer
```

```
create_location general "State College"
```

```
create_location general "Admin building"
```

```
create_location general Employ
```

```
create_location general "First floor" "Second floor"
```

```
set_location Employ "Admin building"  
set_location "Second floor" "Admin building"  
set_location "Admin building" "State College"
```

Assign the node IP03 to location Employ.

```
set_location IP03 Employ
```

Assign the server arch.univ.com to location Second floor.

```
set_location arch.univ.com "Second floor"
```

```
quit
```

To verify that the location hierarchy was created, use any of the methods described in “Step 3: Verify that the information was added to the database” on page 17.

Using an External Customer Database

As mentioned earlier, you should keep the majority of your customer data in an external database. This allows not only for more efficient operation of NNM's object database, but also allows you to keep customer information for which the object database does not provide fields, such as the other information shown in Table 1-1, including key contact name, e-mail address, pager number, or telephone number.

Customer Views provides menu bar options to access such customer information. The menu option expects to launch commands that will access an external database to show whatever information you want your operators to see about your customers.

Customer Views includes a Perl script, `$OV_BIN/ovorgproperties.ovpl`, (*install_dir*\bin on Windows operating systems) which will look for a user-supplied Perl script `$OV_BIN/orgproperties.ovpl`. Customer Views does not include the second script; you will need to create it, using your own customer information.

An example of a suitable `orgproperties.ovpl` is shipped in the `$OV_CONTRIB/CustomerViews` directory (*install_dir*\contribCustomerViews on Windows operating systems). This example script accesses a pseudo-database which is simply a flat, ASCII file, an example of which is also shipped in the same directory.

You will need to modify the script to look in your own database. You could have the script directly access your SQL database, or you could have an intermediate flat file similar to the example file shipped. Your script must be a Perl script called `orgproperties.ovpl` and must be placed in `$OV_BIN` directory (*install_dir*\bin on Windows operating systems).

When a user selects a customer and requests customer properties through the Configuration:Organizations->Customer Properties menu item, the script will access your customer database and view the information for the selected customer.

For More Information

The `ovcustomer` command has numerous parameters that allow you to set other database fields, query the database for specific information about a customer, and print information from the database. In addition, there are ‘shorthand’ versions of the `ovcustomer` commands that simplify its use. These parameters and shorthand commands are described in the UNIX manpage (Windows reference page).

In addition, the online Help on the menu item

`Configuration:Organizations` include more information on using the customer data, as do the `Tasks` in the online help.

2 **Using Customer Views on Management Consoles and in Distributed Environments**

This chapter is not an introduction to Customer Views for NNM. Readers are assumed to be conversant with Customer Views and NNM's key features for network monitoring, data collection, event handling, and so on. Readers who are looking for that information should refer to *Managing Your Network with HP OpenView Network Node Manager* instead. That manual also covers basic configuration of NNM station options that are not exclusive to distributed management. For detailed information on running NNM in a distributed environment you should refer to *A Guide to Scalability and Distribution for HP OpenView Network Node Manager*.

This chapter contains information about running Customer Views on a management console and using Customer Views in a distributed environment.

- Licensing and installation for a management console.
- Configuring management consoles.
- Running in a distributed environment.

All of the documentation for NNM and Customer Views for NNM is written with the assumption that you have activated environment variables on UNIX systems. For information about environment variables, see the reference page in NNM online help (manpage on UNIX systems) for *ov.envvars*.

NOTE

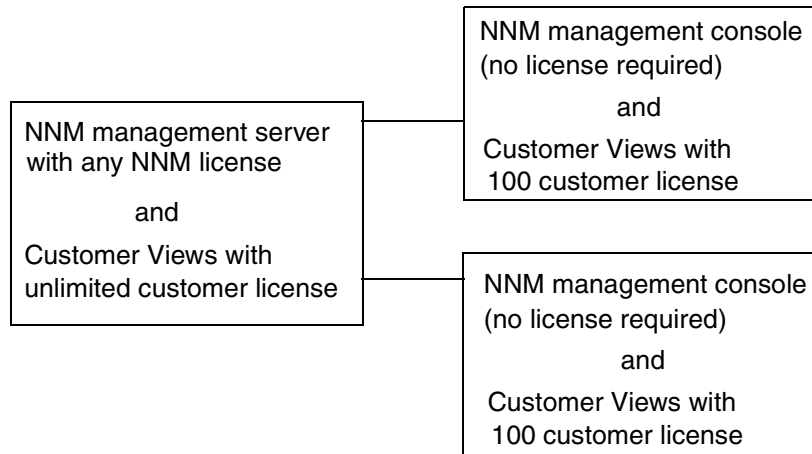
Customer Views supports only HP-UX management consoles. Windows management consoles are not supported. Refer to the Network Node Manager Release Notes, System Configurations topic for the operating systems that are supported with this release.

Using Customer Views on a Management Console

Licensing

To run Customer Views in a management console configuration, the management server and each management console must have a unique Customer Views license and copy of Customer Views installed on it.

Customer Views running on a management console will default to the Customer View license that is installed on the management server that it is connected to. For example, if you have an unlimited customer license installed on the management server and a 100 customer license installed on each management console, Customer Views will use the unlimited customer license.



For information about obtaining a license for Customer Views for NNM, see the *Installation Guide for Network Node Manager and Customer Views for NNM*.

Installation

Installing Network Node Manager on a management console is documented in *A Guide to Scalability and Distribution*. Refer to the Procedures for Scaling NNM chapter.

NOTE

Customer Views supports only HP-UX management consoles. Windows management consoles are not supported.

Follow the basic steps in the manual *A Guide to Scalability and Distribution* and install Customer Views after you install NNM. Remember, you need additional licenses for each management console that you install Customer Views on.

Using Customer Views in a Distributed Environment

Please see the manual *A Guide to Scalability and Distribution for HP OpenView Network Node Manager* a complete understanding of running Network Node Manager in a distributed environment. The information presented in this section does not attempt to explain the distributed environment, it simply presents additional information about running Customer Views for NNM in a distributed environment.

To use Customer Views in a distributed environment, you need to decide where the management of the customer resources will occur. This section discusses the following scenerios.

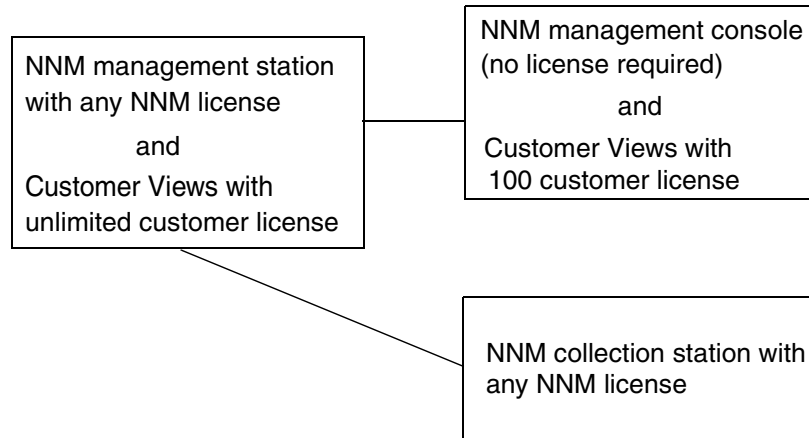
- All customer management is done at the management station.
- The customer management is distributed between the management station and the collection stations.

Managing Customers from the Management Station

When you manage all customer's nodes and devices at the management station, Customer Views needs to be installed on the management station only. The collection stations do not need copies of Customer Views.

When collection stations discover customer's nodes and devices, NNM should automatically be configured to pass the status of these nodes and devices up to the management station, as long as filters that would

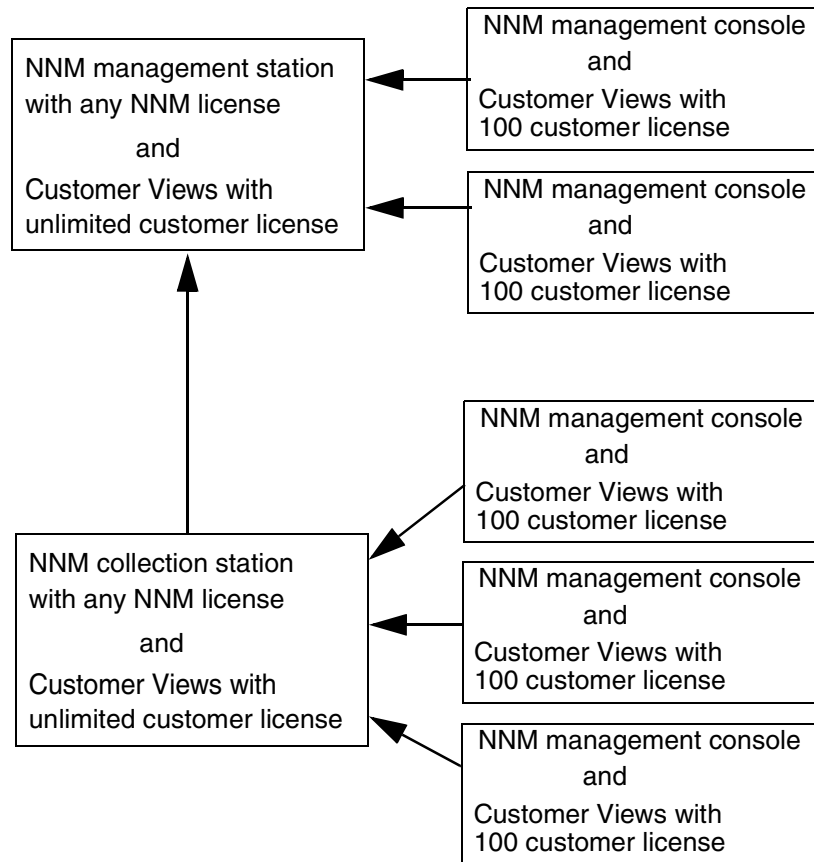
change this configuration do not exist in NNM. This way, Customer Views on the management station can associate the status of the devices and nodes with the customers.



Managing Customers from Both the Management Station and Collection Stations

When you are managing customer's nodes and devices from both the management station and collection stations, you need Customer Views installed on both the management station and each collection station that is performing customer management. This is also true when the collection stations are managing the customer resources locally and may not be passing the status of the nodes to the management station. The collection stations need their own copy of Customer Views to make the associations.

Using Customer Views on a Management Console



You can create associations between organizations and network resources on the management station and collection station. If these organization associations change on the collection station, Customer Views does not automatically propagate these changes to the management station. To view these changes at the management station, you must manually make these changes at the management station.

When a collection station is the primary station for monitoring a given set of network devices, the status of those devices will be updated at the management station during the scheduled polling interval of the primary collection station.

Best Practices

When working in an environment with a distribution model, HP recommends that you use scripts or batch files with the `ovcustomer` utility to configure Customer Views. This will greatly reduce administration costs as well as reduce possible errors for management and collection station synchronization. See the *Configuring Customer-specific Network Management* online manual for more information.

For each management server you should only run one `ovcustomer` utility at any given time. This is to avoid “who wrote last” scenarios. HP recommends that only one management console or the computer console be designated to use the `ovcustomer` utility.

Using Customer Views on Management Consoles and in Distributed Environments

Using Customer Views on a Management Console