The Use and Tuning of the
System Configuration Repository
to Manage System Change

Mark E. Mayotte,
Robin "Stormy" Peters
and
James Metcalfe

Hewlett-Packard
HP-UX Development Laboratory, MS99
3404 East Harmony Road
Fort Collins, Colorado 80528

(970) 898-4141
mark_mayotte@hp.com

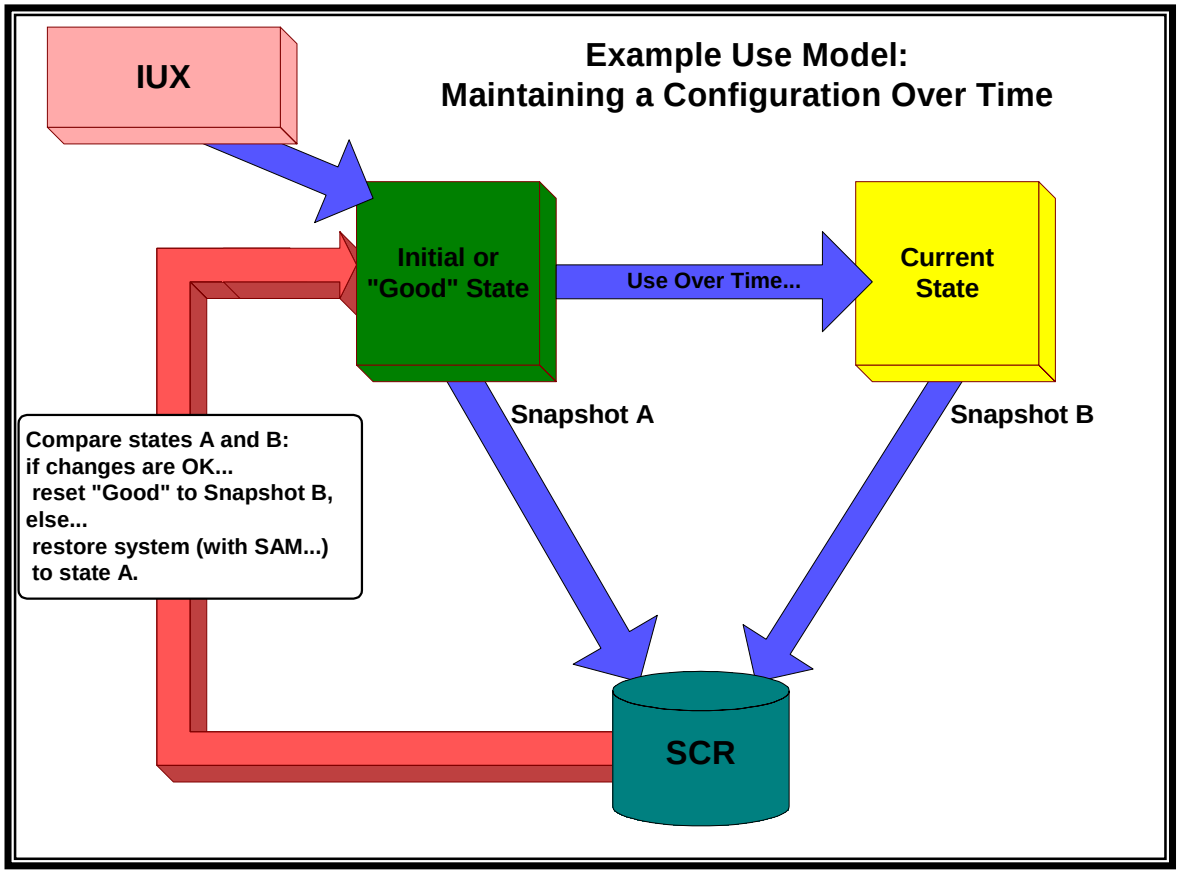

The System Configuration Repository (SCR) is a free tool from Hewlett-Packard that permits an administrator to collect information about a system into a central database as needed. This information can then be used to help manage that system.

In the example above, SCR is used to compare the system's state to a previous, known good state. The administrator can determine if the current state is OK, in which case Snapshot "B" can be labeled as "Good" and used for future reference. If some of the differences are not desired, utilities available on the system (such as SAM) could be used to change the system back to a desired state. SCR would then be used to capture a new "Good" snapshot to serve as the baseline.

As you can see from this simple example, SCR's basic capabilities help to address many common system administration needs.

1 Topics to be covered

We will present a general overview of SCR and several use models that we have encountered when designing and reviewing the tool's current and potential capabilities with customers. This will include practical examples showing use of the SCR commands to accomplish each use model. We will also direct you to more information on the underlying DMI (the Desktop Management Interface) software on which SCR is built. This helps the advanced administrator unlock even more of SCR's potential. We'll conclude with information on how to obtain SCR and DMI, an idea of our future plans, and a summary of how we have responded to customer feedback. The last section is perhaps the most important - without continued feedback and new ideas from our customers, we are not able to produce products that have a chance to meet your needs.

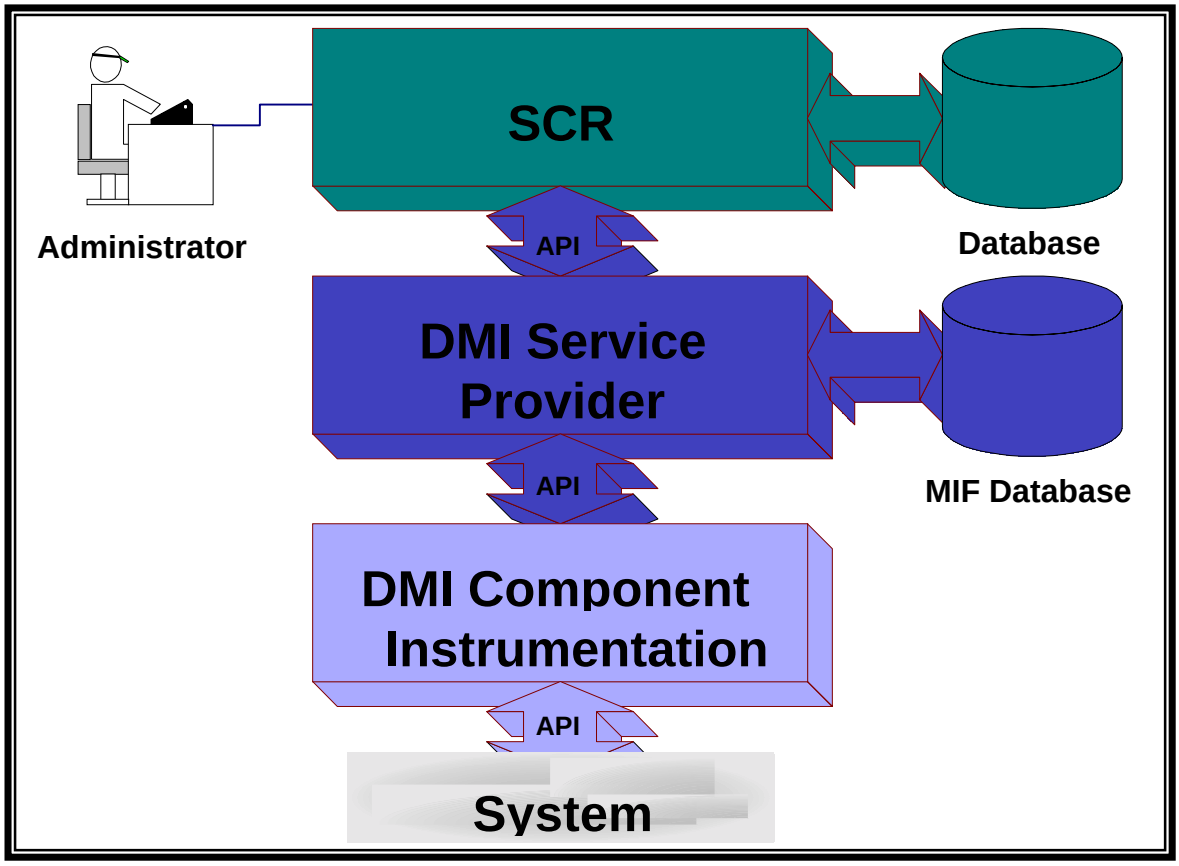

2 SCR Definition and Functional Overview

2.1 Basic Architecture

SCR is a DMI management application. That means SCR gets all the information that it stores from standard API's that are supported by HP's DMI Service Provider. The DMI service provider can be thought of as SCR's collection agent on a managed node.

SCR has the ability to access all of the DMI information available about a machine and store it as a "snapshot" in a central database. SCR follows the client-server model: a single server runs SCR and uses DMI to access client nodes. Exactly what information is captured can be controlled, down to the level of each individual attribute, through the use of "filters."

2.2 Functional Overview

SCR provides access to its functionality through a set of commands:

- scrconfig - adds and deletes nodes in SCR's domain, sets automatic collection and delete schedules, …
- scrstatus - shows completed, scheduled and active snapshot information for each node.
- scrupdate - takes a snapshot of a single node or all nodes known to SCR.
- scrhist - shows all of the snapshots stored in the central database for all nodes.
- scrtag - allows users to name specific snapshots (tagged snapshots are not automatically removed)
- scrviewer - views the contents of a snapshot.
- scrdiff - compares the contents of two snapshots and shows the differences.
- scrfilter - creates, deletes and modifies filters that are used by scrupdate, scrviewer and scrdiff.

The best way to see how SCR works is by example. The following are examples (based on the most recent version of SCR) of how to accomplish a desired task that have been taken from customer interactions…

2.3 Preventative Maintenance and Consistency Management

System administrators can use SCR to give advance warning of any potential problems by monitoring the changes on a set of similarly configured systems.

To use SCR to monitor a system, the system administrator would:

Initial Set-Up Steps NOTE: The first 7 steps below are very similar for each scenario. They will be referred to as the "initial set-up steps" in later examples.

First install and configure DMI on all of the systems that are being watched.

Install and configure SCR and DMI on the central management system (This system is not necessarily one being managed).

From the central management station, run:

`scrfilter -o Probe my_probe`
Dump the definition of the *Probe* filter to the file named "my_probe" to permit customization. See below for more information on the importance of the *Probe* filter.

`vi my_probe`
Tune the *Probe* filter by editing the file.

`scrfilter -i Probe my_probe`
Load the customized version from the "my_probe" file into SCR replacing the previous version of the *Probe* filter.

`scrconfig –n +node1 +node2`
Add systems.

`scrupdate -a`
Take an initial snapshot of the systems as they are now. You may want to take a look at what is captured (scrviewer) and further tune the Probe filter until you are capturing the right amount of data. You can use the scrstatus command to see then a collection is completed.

`scrdiff node1:latest node2:latest`
Ensure that they are consistent from the outset.

`sam; scrupdate; scrdiff…`
Make any changes needed, take another snapshot and compare until good.

`scrconfig –i 1w`
Set SCR to collect data weekly.

Periodically (perhaps with a script), check to see if the nodes are still consistent. From the central management server, run:

`scrstatus`
Ensure that SCR successfully collected data.

```
TIME      (START - STOP)       NODE                STATUS    DETAIL
02/21/2000 23:30 – 23:34 MST   node1               Completed
02/27/2000 23:30 MST           node1               Scheduled
02/21/2000 23:34 – 23:38 MST   node2               Completed
02/27/2000 23:30 MST           node2               Scheduled
```

```
#  Scrhist
```
See what data is in the database - SCR will not store a snapshot if there are no differences from the previous snapshot. SCR automatically tags the first and last snapshots with "oldest" and "latest" respectively.

```
NODE             TIME                   ERR  TAG
node1            05:20 02/07/2000 PST       oldest
                 23:37 02/14/2000 PST
                 23:34 02/21/2000 PDT       latest
node2            05:23 02/07/2000 PST       oldest
                 23:42 02/14/2000 PST
                 23:38 02/21/2000 PDT       latest
#  scrdiff node1:2000022337 node1:latest
   scrdiff node2:2000022342 node2:latest
   scrdiff node1:latest node2:latest
   …
```

See what has changed since the last time data was collected. You would check for changes on each node over time, and between each node.

2.4 Troubleshooting

System administrators can use SCR during troubleshooting scenarios. They can use SCR to view the state of their system and compare it to previous known states. In this way, they can see what changes (authorized or not) have happened on the system.

In order to set SCR up for use in a troubleshooting scenario, the system administrator would:

```
#  Complete the initial set-up steps.
#  scrtag -a node1 golden 200002071205
```
Label or “tag” the system to record this state. It is assumed that the current system configuration is "golden" for this example.

```
Tag "node1:golden" is registered for 02/07/2000 12:05 PST.
#  scrtag -a node2 golden 200002071205
Tag "node2:golden" is registered for 02/07/2000 12:05 PST.
#  scrconfig –i 1d
```
Set SCR to collect data once a day or to your desired interval. Choose a frequency that matches how often your system's configuration is changed.

When troubleshooting, the system administrator would run:

```
#  scrupdate –n node1
```
Collect current data from the system having problems. You can skip this step if an automatic snapshot was taken after the problem surfaced.

```
#  scrdiff node1:latest node1:golden
```
Compare the current data to a known good state to see differences. If the system administrator knows, for example, it’s a networking problem, just view the relevant data:

```
#  scrdiff -f Network node1:latest node1:golden
```

2.5 Inventory/Reporting

System administrators can use SCR to generate a document that says “this is what my system looks like”. It can be used to maintain inventory, to plan for future changes or to help third party software vendors establish what you need.

In order to set SCR up for use in a reporting or inventory scenario, the system administrator would:

Complete the initial set-up steps.

`scrviewer node1:latest | lp –dprinter`
Generate a full listing of each system's configuration from the central management station. Print the data from node1, or use the data from scrviewer to fill in your standard report, or save the output in a file to save/format as needed.

SCR's output is not formatted nicely; it reports information in its DMI structure (more on that below). This is not always the most usable for an end-user report. However, it is reported in attribute/value pairs that facilitate processing in scripts to generate reports.

2.6 Change Management.

In many shops, changing systems needs to be accomplished in a careful, often centralized, controlled fashion. SCR can be used to complement a centralized or decentralized change management process. System administrators can use SCR to notice and document changes on a system because SCR:

- Takes periodic or on-demand snapshots.
- Compares the current state to previous states (in any combination desired).
- Keeps a record of all changes.

In order to set SCR up to use as a change management agent, the system administrator would:

Complete the initial set-up steps.

`scrupdate -a`
Take an initial snapshot of the systems as they are now. This provides a base from which to detect future changes.

`scrconfig –i 1w`
Set SCR to collect data weekly (or more/less after depending on your environment).

`scrhist`
On the same interval as the collections, email a history report to the system administrator. Remember that SCR will not store a snapshot if there were no changes. The report below was generated on the 22nd and shows that node1 had changes each week while node2 has not changed between the 14th and the 21st.

```
NODE            TIME                    ERR  TAG
node1           05:20 02/07/2000 PST       oldest
                12:05 02/14/2000 PST
                12:05 02/21/2000 PDT       latest
node2           05:23 02/07/2000 PST       oldest
                12:10 02/14/2000 PDT       latest
```

`scrdiff node1:200002141205 node1:latest`
After receiving the email, the system administrator runs a difference report from the central management system to see what the changes were. No report would be run on node2 as there were no changes.

2.7 Disaster Recovery

System administrators can use SCR to check that a system has been fully restored to its previous state.

To use SCR to in a disaster recover scenario, the system administrator would:

Complete the initial set-up steps.

`scrconfig –i 1w`
Set SCR to collect data daily, or less often if your systems are more stable.

In the case of a disaster recovery situation, restore the system (with Ignite UX or other software). Make sure DMI is installed on the system. (And configured to allow the central management system access.)

`scrupdate –n node1`
Collect data from the restored system.

`scrstatus`
Make sure the collection was successful.

```
TIME      (START - STOP)       NODE                STATUS    DETAIL
01/07/2000 09:23 – 09:27 MST   node1               Completed
```

`scrhist`
See what snapshots are in the database.

```
NODE             TIME                    ERR  TAG
node1            12:20 01/05/2000 PST       oldest
                 12:05 01/06/2000 PST
                 09:27 01/07/2000 PDT       latest
```

`scrdiff node1:200001061205 node1:latest`
Compare the current data with a set of data that existed before the recovery to discover differences.

By now you see how SCR can be used in many different ways. The disaster recovery scenario can be extended to include non-disaster situations such as preparing for a reboot and making sure all is well afterwards. It can also be used to document the configuration prior to an OS upgrade and then checking for changes afterwards.

2.8 What it collects and stores, and how you view it

In order to best understand how to manage systems using SCR, it is necessary to understand how SCR stores data. Put briefly… it stores it in DMI format. DMI provides a three-level structure for data: *Components, Groups* and *Attributes*. SCR mirrors that format as can be seen in the following example data:

```
COMPONENT NAME                                  VALUE
    GROUP NAME
        ATTRIBUTE NAME
"HP-UX Standard Groups Definition"
    "Host Storage"

        [Host Storage Index]                    1
        [Storage Type]                          7:CompactDisc
        [Description]                           " TOSHIBA   CD-ROM XM-5401TA"
        [Allocation Unit Size]                  0
        [Total Allocation Units]                0
        [Allocation Units Used]                 0
        [Storage Allocation Failures]           0

        [Host Storage Index]                    2
        [Storage Type]                          4:FixedDisk
        [Description]                           " SEAGATE   ST32151N        "
        [Allocation Unit Size]                  1024
        [Total Allocation Units]                2097342
        [Allocation Units Used]                 0
        [Storage Allocation Failures]           0

        [Host Storage Index]                    3
        [Storage Type]                          4:FixedDisk
        [Description]                           " SEAGATE   ST32430N        "
        [Allocation Unit Size]                  1024
        [Total Allocation Units]                2097342
        [Allocation Units Used]                 0
        [Storage Allocation Failures]           0
```

Each data item is identified by its place in the hierarchy (indicated by the indentation level). Data attributes are always shown in attribute/value pairs. The attribute name is contained between "[" and "]" delimiters. The value follows the attribute name. If there is more than one device, the attributes for each are shown in order (this is shown in the above example with 3 storage devices).

The difference report uses the same basic structure with two columns used to show the two values being compared. In the following example the system administrator has changed the DMI information that documents the primary user of the system. The changes are highlighted with an "*" in the first column.

```
  COMPONENT NAME                          BASE              TARGET
      GROUP NAME
          ATTRIBUTE NAME
  "HP-UX Standard Groups Definition"
      "General Information"
          [System Name]                   node1             node1
          [System Location]               BLDG1             BLDG1
*         [System Primary User Name]      "Joe User"        "Jane User"
*         [System Primary User Phone]     555-1234          555-5678
```

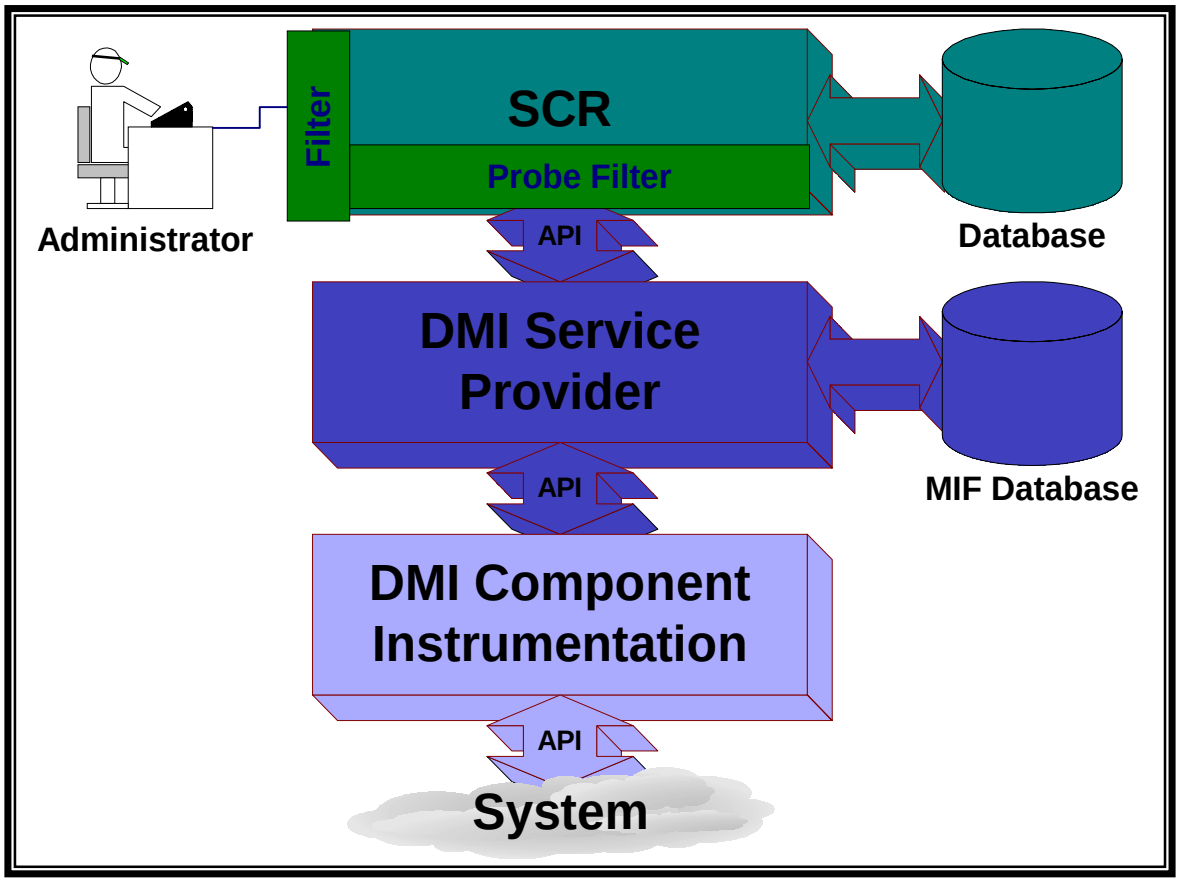

2.9 Filters

2.9.1 Role of Probe Filter

The Probe filter determines the information that is taken from all of the available DMI instrumentation. The Probe filter determines what is captured in SCR's database. It is very important to set the Probe filter carefully. Some points to remember are:

- Setting the filter too "loose" captures too much data and takes a lot of time to collect. Setting the filter too tight can cause you to miss data that might be important to your installation.
- The Probe filter setting is a large factor in determining the number of nodes that can be managed from a single server. It determines the database size and the time required to capture data from each node. On a mid-range server scrupdate takes about 3-5 minutes to run with the default Probe filter setting. On slower system a snapshot takes about 10 minutes. A full snapshot (all aspects of the Probe filter set to "On") will take more than an hour and occupy more than 8 Mbytes of space.
- In order to focus attention on configuration changes, SCR will not store a snapshot in the database if it is the same as the previous snapshot. Certain DMI values will always change (such as process information and used disk space) and if included in the Probe filter will cause SCR to always save the snapshot in the database. The default version of the Probe filter is set to ignore these attributes. You may want to turn them on for your site.
- The default Probe filter is pretty loose in that it errs on the side of capturing data at the expense of grabbing more than each site probably wants or needs.

It is not possible to correctly anticipate the different needs of our customers which is why we have provided the Probe filter capability. The Probe filter allows the customer to tune the data collected by SCR to meet their needs. It is important that you play with the Probe filter until you are comfortable with the amount of data SCR collects about your system. After you have a Probe filter definition that works for you… then you can deploy SCR into your production environment.

If you change the Probe filter after some snapshots have been collected, you will see the difference with the scrdiff command. SCR displays missing data with a "-." In the following example the Target snapshot did not include the General Information group:

```
      COMPONENT NAME                              BASE              TARGET
          GROUP NAME
              ATTRIBUTE NAME
      "HP-UX Standard Groups Definition"
          "General Information"
    *         [System Name]                       node1               -
    *         [System Location]                   BLDG1               -
    *         [System Primary User Name]          "Joe User"          -
    *         [System Primary User Phone]         555-1234            -
```

2.9.2 Display Filters

In addition to the Probe filter, SCR supports display filters. These filters are used to control the output of the scrviewer and scrdiff commands to a desired subset of the all of the information collected by the Probe filter. We include a set of these with SCR. These can be seen with the scrfilter command:

```
#  scrfilter -l
   Disk
   FileSystem
   LVM
   Network
   Patch
   Probe
   Software
   SystemProperty
   Template
```

The "Template" filter is provided as a starting point to create a new filter. You can use the "-o" and "-I" options of scrfilter to dump a filter definition to a file and load it back in when you are done editing. You can also directly set the value for given attribute from the scrfilter command line. Here is the Disk filter (the actual output is quite long as all attributes, whether set to On of Off, are included. Only the groups with attributes set to "On" are shown below):

```
#  scrfilter -o Disk disk_filter
#  cat disk_filter
   [Filter File Version]: 1
   [Filter Name]:  Disk
   [Data Version]: 1
   [Extension]:    off

   Component:      "HP-UX Standard Groups Definition"        on
     Group:        "Host Storage"                            on
       Attribute:  "Host Storage Index"                      on
       Attribute:  "Storage Type"                            on
       Attribute:  Description                               on
       Attribute:  "Allocation Unit Size"                    on
       Attribute:  "Total Allocation Units"                  on
       Attribute:  "Allocation Units Used"                   on
       Attribute:  "Storage Allocation Failures"             on
```

The first four fields are control values for SCR… don't mess with those ☺.

3 Some General Information

3.1 DMI and Product Availability

As was mentioned above, DMI can be thought of as SCR's collection agent that runs on each managed node. If you want to add information to SCR beyond what is available from HP, you do so by adding DMI instrumentation.

DMI is an instrumentation standard that HP provides for the 10.20 and 11.0 versions of HP-UX. We presented a paper on DMI describing how to create your own instrumentation at InterWorks 1999. You can also find out a lot more by referencing HP's Software Depot web site at http://www.software.hp.com/products/DMI/ (this is the same site where SCR can be found). DMI and SCR are also available on the Application Release media as:
B6814AA for Series 700 10.20, B6260AA for Series 800 10.20, and B6816AA for Series 700 & 800 11.0.

The website contains a complete DMI Developers Guide that documents all of the supported API's for programmers. A DMI User's Guide is also provided that documents all of the DMI instrumentation that HP provides as part of the DMI product. If you want to understand just what those attributes are that SCR is reporting, this is where you look.

3.2 SCR and DMI Future Plans

We first released SCR version 1.0 in December 1999. We provided an update of SCR in March 2000. The next release of SCR is planned for the end of 2000 or early 2001. We are working on a set of phased releases that permits us to provide the SCR product quickly. We then provide incremental capabilities to meet our customer's needs. Reference the Appendix for a summary of some of the customer feedback we have gathered and what our plans are for addressing these needs. Two of the higher priority enhancements include:

3.2.1 A Full Function GUI

The current SCR product provides a great deal of capability through a command line interface. This provides great flexibility for general use and for writing scripts to automate routine tasks. We are investigating providing a GUI that would permit the browsing of snapshots and their contents, viewing and changing SCR control information, and possibly integrating in DMI so you have access to stored snapshots and dynamically determined values side by side.

3.2.2 Reporting

As you can see from the above there is a need to provide some report formatting capabilities to make the information in SCR easier to digest. We are looking into providing filter/format pairs that would help address this. A filter/format pair is an SCR filter that restricts the information to a specific area and then "pipes" the information into a formatting script that takes care of the rest. The attribute/value format of SCR output makes it relatively easy for the end user to construct customer reports, but it would be valuable to provide a set as part of the shipped product.

4 Acknowledgements

First, we'd like to acknowledge the many customers that took the time to sit down with us, either at their site when we were able to visit, or at last year's InterWorks conference. Without these inputs we could not have come up with a product designed for our customers (and, no, we don't think we are done yet).

SCR was a collaborative project between Hitachi and HP. Hitachi has been working closely with HP to enhance the HP-UX environment for the benefits of both of our customers. Hitachi has enhanced our manageability offering from the HP-UX Development Lab with the EMS Kernel Resource Monitors, enhancements to the DMI instrumentation shipped as part of DMI, and now SCR.

5 Appendix - Summary of Customer Feedback

5.1 Customer Feedback

SCR is a free tool from Hewlett-Packard created in response to customer requests. During conversations with customers, we discovered that many people have created their own scripts and applications to collect inventory data on HP-UX systems. The current implementation of SCR is a direct response to input we have received from customers both at customer visits and from previous InterWorks conferences. In addition, we are constantly refining SCR and improving future versions in response to customer feedback.

This table is provided to document the feedback we have received and how we have responded to that feedback. If you look closely, you'll see that we have implemented many of the requests! Please continue to help us evolve SCR to meet your needs - let us know what those needs are.

Feedback	Source	HP Response
More flexibility is setting data. Customers said that the amount of data being collected and stored was more than they need.	Customer Visits	The default settings for collection have been greatly pared down, but are still fairly open so desired data is not lost. Users who want more data can change the default settings to include the data they are interested in.
Merge local with remote. Customers would like to have SCR running on each system. If there are changes on the system, then they would like to update the SCR database on the server. This was seen as a way to reduce network traffic and reduce the size of the central database.	Customer Visits	SCR does not merge data, but it only stores a snapshot if it detects differences. The checking is done from the server.
Reports, specifically html formatted reports.	InterWorks 1999, Customer Visits	We are still investigating what data and what formats should be used for reporting. This will be added to a future release of SCR as it is one of the most commonly heard requests.
Predefined filters.	Customer Visits	SCR ships with seven predefined filters (in addition to the Probe and Template filters): Disk, FileSystem, LVM, Network, Patch, Probe, Software and SystemProperty. We will add others as customer feedback suggests.
Snapshots that never get deleted. Users would like to have some data that never "expires" and is therefore never automatically removed from the database.	InterWorks 1999, Customer Visits	Beginning with SCR 1.0, any snapshot that has a user defined tag will never be automatically removed from the database.
Customers want it to be easy to pick two nodes and compare latest snapshots or pick one node and compare either the latest two snapshots or the current data against the last snapshot.	Customer Visits	SCR provides a simple GUI as part of the Service Control Manager. In the Service Control Manager, if you pick two nodes and select the "Compare Inventories" tool, the default will be to compare the latest snapshots. If you pick one node, and select the "View Inventory Changes" tool, the default will be the two latest snapshots on that node.

Feedback	Source	HP Response
Reconfirmed that customers want to be able to see SCR data from within HP's ITO and ITA products.	InterWorks 1999, Customer Visits	SCR has been integrated with the Service Control Manager. Merging with other applications is under investigation.
From the GUIs, users want to be able to print what they are looking at (or save to a file that can be printed). They want the option to be in the view dialog.	InterWorks 1999	As part of the integration with the Service Control manager, "Save to File" will be available from every dialog that displays inventory data. This capability will be preserved in future GUI's.
Users would like to collect additional data, other than the DMI data provided by HP, such as the contents of files, which IUX image was used to ignite the system, serial numbers , system owner, 3rd party software, etc.?	Customer Visits	DMI provides the ability to extend SCR. More instrumentation from HP as part of the shipped product is currently under investigation.
Users would like to be able to filter on the value of the data.	Customer Visits	User access to SCR's internal database APIs is under consideration for a future release. When the APIs are available, users will be able to filter on the value of the data.
Users would like SCR agents on other UNIX and PC platforms.	InterWorks 1999	This is currently under investigation.
Customers would like scrdiff to diff patches and update "older" systems.	InterWorks 1999	Currently users can diff just patches using the patch filter. Currently, SCR does not change any values on the system.
Users would like access to database APIs to directly manipulate inventory data.	InterWorks 1999	Access to SCR's internal database APIs is under consideration for a future version of SCR.
Capture installed patches	InterWorks 1999	Not only does SCR include the ability to capture patches, but a patch filter is included. Using this filter, users can view just the patches on a system. (On 10.20, they will see all of the software installed, which includes patches).
Can you configure SCR to take a snapshot every time a system boots?	InterWorks 1999	SCR does not include this capability. You can, however, add a call to scrupdate as part of the boot sequence.
Is there a way to insert a baseline? One that never expires and can be manually edited?	InterWorks 1999	With SCR users can create a baseline, by giving your baseline snapshot a "tag" so that it will not be deleted. Currently users cannot manually edit that snapshot.
How many snapshots can you do in parallel?	InterWorks 1999	Currently, SCR captures data serially and we recommend a limit of 30-50 systems per Central Management Server. The exact number that can be supported depends on your configuration, machine performance, and the setting of the Probe filter. The next major release of SCR is investigating the ability to capture data in parallel.
Can it do a PC inventory?	InterWorks 1999	We are currently investigating capturing information from PCs through DMI's interoperability capabilities.
Does it take corrective actions based on differences found?	InterWorks 1999	Currently, SCR does not change any values on the system.