

Slightly misc notes on HP Service Manager

Greg Baker (gregb@ifost.org.au)

January 18, 2012

Contents

1 Short-cut keys for the web interface	3
2 Tasks in a roll-out plan	3
3 Suggestions on populating CI database	5
4 Restarting processes automatically	6
5 What the processes do	7
6 TRCLIENT	7
7 Doing development	8
7.1 Assemble a lot of systems	8
7.2 Unload script for forms and basic configuration	8
7.3 Unload script for forms and basic configuration	10
8 Troubleshooting Access Issues for Knowledge Documents	11

9 What we use CIs for	11
9.1 CI Reporting	12
10 Giving releases a different numbering scheme	13
10.1 Create relmgmt numbering record	13
10.2 Alter the cm3r format control	14
11 Locking and the linker	14
11.1 The record has been changed since you last updated it . .	14
11.2 When the linker never updates	15
12 Installation Notes	15
12.1 Installing the windows interface on Windows 7	15
12.2 Installing the Windows interface on Linux	15
13 Operator Records	15
14 Installing the Web tier on ServiceManager 9.2 on Linux	16
15 DVD tricks	16
16 What to think about when bringing a new team or Service Desk to a ServiceManager environment	17
17 Last-fortnight Training Requirements	18

List of Tables

2 Web interface key bindings for 7.X	4
--	---

3 Standard Training Matrix 20

List of Figures

This document is where I record observations about HP's ServiceManager product which I haven't seen documented coherently elsewhere, or which I thought were interesting.

1 Short-cut keys for the web interface

Table 2 lists the keybindings for the web interface (which HP recommends using now, instead of the windows interface).

2 Tasks in a roll-out plan

I haven't documented the necessary steps to start using change management sensibly, nor problem management, but I thought figure 2 might be useful to some people anyway.

I did include one very important set of steps which isn't *strictly* necessary, but which can make the difference between a mediocre Service Manager implementation and a fabulous, company-transforming one. That is – first figure out what the real priorities are for

In most organisations there is a single step in a single workflow which is the “bottleneck” – the process which is limiting the organisations ability to make money (or do whatever it is the organisation is supposed to do).

Sometimes there may be more than one, and sometimes they can be present across a couple of workflows, but in general it is possible to identify a very small number of the worst problems.

These are the places that need to have priority on IT – if work is slowed down on a bottleneck because of IT problems, the organisation can never catch up on it. IT problems in non-bottleneck areas can be deferred for a short period of time, because the work can usually be caught up somehow.

The default Service Manager logic for priority is simply the average of

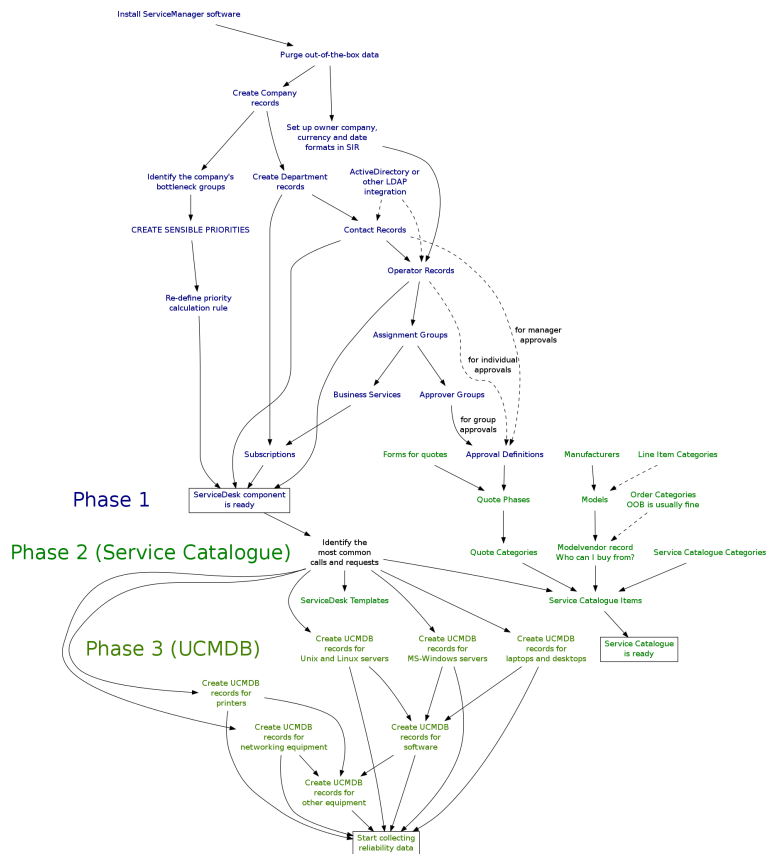
Keystroke	Binding	notes
Alt-c	SpellCheck	Opens favourites in IE by default. Can use Ctrl-Alt-c instead.
Alt-g	Messages	
Alt-i	Help	Doesn't seem to work in IE.
Alt-n		Collapse or Expand Navigation panel
Alt-o	About	Opens tools in IE by default. Can use Ctrl-Alt-o instead.
Alt-p	Print	
Alt-r	Refresh	Refreshes page in IE by default. Can use Ctrl-Alt-r instead.
Alt-1		Manage Favourites
Alt-2		Add to Favourites
F1	Help	Disabled by default, but works in IE anyway.
F2	Magnify	View field magnified
F4	Help	But doesn't work in IE.
Shift-F1	Save	Only works on Mozilla.
Alt-F1	New	Can also be Clone or Re-open
Alt-F2	OK / Refresh	
Alt-F3	Back	
Alt-F5	Delete	Can also be "Undo"
Alt-F6	Search	
Alt-F7	Clear	
Alt-F8	Count	Also can be "Find related record"
Alt-F9	Fill field	
Alt-F10	Next	
Alt-F11	Previous	
Alt-F12	Clocks	
Alt-F22	Apply Template	
Ctrl-F4	Save	But this closes the window in IE.
Ctrl-F12		Transfer focus from left navigation pane to main viewing pane.

This document was generated from web.xml in a default SM 7.10 installation, the accessibility information and mouse-overs in the web interface.

Table 2: Web interface key bindings for 7.X

the affected group¹ and the urgency². This means that a low priority company-wide issue is given priority over a critical problem on a bottleneck workflow!

How to install and configure HP Service Manager (c) Greg Baker (gregb@ifost.org.au)



3 Suggestions on populating CI database

Step 1 is to create the business services definitions – otherwise service desk calls can't be created! But that is all that strictly must be done in order to roll out ServiceManager.

There can be some value in populating the CI database further. Here is my checklist of CIs to look for – first decide what is going to be valuable to maintain in ServiceManager, and then perhaps use some of the suggested techniques for populating it.

¹ 1 = Enterprise, 2 = Site, 3 = Multiple Users, 4 = Individual User

² 1 = Critical, 2 = High, 3 = Average, 4 = Low

Windows servers Often visible in ActiveDirectory, but watch out for machines which are not part of a domain.

Windows desktops ActiveDirectory will usually have a complete list; otherwise asset management may document what laptops and desktops are out there.

Linux servers Look in nagios, cfengine, RedHat Network and so forth. Don't forget firewalls, proxy servers, invisible packet-sniffing machines (e.g. machines running snort or other intrusion detection software).

Unix servers Similarly to Linux servers, but many older legacy Unixes may not appear in any modern administration tools.

Printers Often surprisingly hard to get all of them; especially the non-networked ones. But it's often extremely valuable to getting costing information about service desk calls for printers by make and model to understand what vendors' printers have costly issues; and also to measure printer downtime.

Routers and switches And other network infrastructure – such as VPN concentrators, load balancers and accelerators.

Virtual machines Look for VMWare, Hyper-V, Xen and Virtuozzo on the server side, and VirtualBox, VirtualPC and Parallels on the desktop.

Server software Sometimes it helps to audit per operating system.

Desktop software This is a never-ending list. Only if you have very strict desktop lockdown will you ever maintain this list manually.

Mainframes and minicomputers This is usually not a long list, so it is usually neither particularly expensive nor particularly valuable to keep in ServiceManager. It can be useful for understanding the real costs of legacy systems. Ask for lists of OpenVMS, MPE/ix, MVS, z-series machines and so on.

Other PBX systems; voip phone handsets; server racks; UPS (uninterruptible power supplies); environmental monitors. The very enthusiastic might include patch panels and cooling units.

4 Restarting processes automatically

1. Create an anubis agent record with the form `info.startup.g`

Type anubis

Name anubis

RAD Application scheduler

Class anubis
Wakeup interval 30
Priority 1

2. Create an anubis schedule record with the form `schedule.looksee.g`

Name anubis
class anubis
expiration Some timestamp earlier than now
scheduled class anubis
repeat interval 00:00:30
application apm.anubis

5 What the processes do

anubis See section 4.

lister Runs `apm.server.initer`, which updates any global lists which are out-of-date.

linker Takes tickets from “Open - Idle” state to “Open - Linked” when there is a record related to an interaction.

inactive Logs out sessions which are idle.

TRCLIENT See section 6

6 TRCLIENT

Found on HP Communities:

TRCLIENT is a dedicated server thread which is responsible for servicing jgroups requests.

Jgroups requests are categorized into 3 types:

1. requests which can be handled in java. (Ex: User Session Count)
2. requests which need to be hooked into shared memory. (Ex: status screen to get the list of locks, list of users etc.,)
3. requests which are processed asynchronously. (Ex: IR Invalidation, IR Remove and IR regeneration.)

For requests of type 1 and 2 we will have one TRClient (Dedicated Server Thread)

For requests of type 3 we will have one TRClient (Dedicated Server Thread)

So at any point of time, you should see a max of 2 TRClients per Container and min of 0.

(Thanks to Brav0).

7 Doing development

7.1 Assemble a lot of systems

You will want to have $3 + (\text{thenumberofdevelopers})$ systems – one for production, one for UAT, one for testing and one each for each developer. The developers' systems can be virtual machines running SQL Express (or Oracle Express) and whatever operating system you run production, UAT and testing on. (Unless prod/UAT/testing is on HP-UX or Solaris, in which case the developers could run Linux.)

7.2 Unload script for forms and basic configuration

Create an unload script which takes your production environment's forms, scripts, queue, views and so on. Make sure that when you apply this to a fully purged system that you end up with a working system.

This unload script will need to unload at least:

- Object
- Process
- States
- application
- code
- datadict
- displayevent
- displayoption

- displayscreen
- enclapplication
- eventfilter
- eventmap
- eventregister
- formatctl
- format
- joindefs
- link
- menu
- querystored
- scmessage
- screlconfig
- scripts
- triggers
- validity
- wizard
- ScriptLibrary
- cm3rcategory
- cm3rcatphase
- cm3profile
- cm3tcategory
- cm3tcatphase
- knownerrorcat
- knownerrorphase
- rootcausecat
- rootcausephase
- rootcausetaskcat
- category
- FolderDef

- FolderRights
- contextDefinition
- contextAction
- globallists
- location
- pmenv
- userrole

It will often also include some or all of the following:

- company
- department
- operator
- contact
- group
- slo
- sloavail
- unload

Schedule a regular unload using this script that you have made, and copy the resulting file away from production to somewhere accessible to developers.

Developers will use this to take their development environments to being a clone of production.

7.3 Unload script for forms and basic configuration

Start with a baseline; a system which has all out-of-box data purged.

Don't forget that if you want to load a previously-unloaded script from the Windows/Unix command-line, the `.unl` file has to be in the same directory as the `sm` executable.

On MS-Windows run:

```
sm file.load filename.unl NULL NULL winnt
```

On Unix / Linux run:

```
sm file.load filename.unl NULL NULL unix
```

8 Troubleshooting Access Issues for Knowledge Documents

- Go to the article in question. If necessary, do an advanced search, and type in the document id.
- Hit “Edit” on the document, and then go to the “Categorisation / Flags” tab. Make a note of all the categories the document appears in.
- There are only seven profiles. Most of the time helpdesk staff will be KCS I for the areas they work in. Specialist knowledge management staff will want to be KCS II or III for every area. Default level is usually OK for everybody else for everything.
- Go to Menu Navigation → Knowledge Management → Manage KM Document Categories.
- For each category you noted above, go to that category and select the permissions tab. Note what groups have access at the level you care about (e.g. if the user needs edit access, just look for the KCS I and higher profiles; if they need read access, look for the default groups).
- From the subset of groups, figure out which is most appropriate, and put the user into it. (e.g. from the operator record, or from Menu Navigation → Knowledge Management → Knowledge Groups).

9 What we use CIs for

Managing a CI (Configuration Item) database can be very time consuming, and it can be hard to find a return on investment from it.

Is it possible to have an incomplete CI database and still derive value from it? Quite probably. The trick is to focus on what is going to bring the best returns.

In decreasing order of value to an organisation (usually), the main uses for a CI database are:

1. For reporting, as discussed in section 9.1 on page 12.
2. As an asset register for keeping track of items which will need to be depreciated in the company accounts. Most organisations use AssetManager or AssetCenter for this kind of task, but it can be done in ServiceManager.
3. As a way of uniquely identifying what a Known Error or Change record is associated with.
4. As a filing place for storing documents that you otherwise might have put somewhere else. (e.g. licensing info, vendor management, as-built documentation).

So an incomplete database which at least supports reporting (and includes all the expensive, depreciable assets if ServiceManager is being used for this), it can be populated on-demand when there it is needed for any of the other two uses.

9.1 CI Reporting

Reporting is the function for which a CI database provides the most value.

The most effective reporting is “A vs B” reliability reporting – for example, based on past incidents, interactions, problems and outages, should we use Windows or Linux? Should we buy HP or Dell?

To make this happen you need at least two models for one subtype of CI, many occurrences of that CI, some guess at the cost of an incident, interaction and problem and at least a vague measure of the cost of an outage. Also, if there is a significant price difference between the models (e.g. you are comparing iPhones to Nokias) then don't forget to include that in the report as well.

Note that very often this *does not need the individual CI to be identified*. Plan ahead what “A vs B” reporting you plan to do; the service desk only needs to record “A” or “B” as the CI data for you to do this. For example, if you want to compare laptop vendors, then get the service desk to complete “Affected CI” with “HP Laptops”, “Dell Laptops”, “Apple Laptops” or “Other Laptops”. This doesn't take long (it's far faster than getting a user to give a laptop name or serial number) and you can still derive a lot of value from it.

The other kind of reporting which can be valuable is a trend on outages (and interactions, incidents and problems). Are systems getting worse over time?

Linear trending is OK, but be aware that both software and hardware tend to have U-shaped graphs of issues over time – they start out with many issues, become better and eventually end up with so many issues that they are finally disposed of.

Finally, being able to whinge to a vendor or report to customers on the number of events on each CI can be useful. But again, plan ahead what level of detail you will need in order to report that – it usually won't even need individual computers to be identified.

10 Giving releases a different numbering scheme

Release Management is just another category of Change according to ServiceManager, so it gets numbered the same as any other change.

To fix this, create another numbering record for release management changes, and adjust the `cm3r` format control.

10.1 Create relmgmt numbering record

Go to Menu Navigation → Tailoring → Tailoring Tools → Sequential Numbers

Create a new record with the following parameters:

Class relmgmt

Last Number 1

Length 5 (or whatever you want)

Prefix R

Remember to hit Add.

10.2 Alter the cm3r format control

Type in `fc` in the command-line, or otherwise go to Menu Navigation → Tailoring → Format Controls

Search for the `cm3r` format control record, and select the Subroutines tab. From the Options menu, select Show Expanded Form.

After the existing definitions, add another as follows:

Application Name `getnumb.fc`

Names • `record`

- `name`
- `prompt`

Values • `$file`

- `relmgmt`
- `number`

Before `true`

Initial `(number in $file)~#"R" and category in $file="Release Management"`

Remember to hit Save.

11 Locking and the linker

11.1 The record has been changed since you last updated it

The service desk, incident and problem applications have a checkbox in their profiles called “Lock on display”. Turning this on removes the annoying problem of the record changing while you are working on a ticket (and thus invalidating what you have typed).

In exchange, you end up with the annoying problem of not being able to see a record if someone is looking at it.

If you log out badly, ServiceManager eventually figures this out and frees up the lock.

11.2 When the linker never updates

Tickets can get stuck in the “Open - Idle” state instead of going to “Open - Linked” because the `linker` is unable to update the record. This can happen if the record is locked or otherwise un-modifiable for the 30 times the linker tries. (By default, once every minute).

12 Installation Notes

12.1 Installing the windows interface on Windows 7

If you try to install the Service Manager Windows 7.X client on Windows 7, it errors out. It has something to do with the installer GUI because it works fine if you run `setup.exe /silent`.

12.2 Installing the Windows interface on Linux

Actually, it just works. I used CrossOver and told it to install from `setup.exe`.

13 Operator Records

If you have a large number of assignment groups, and staff who move in and out of assignment groups regularly, the Assignment Groups input area in the operator record can get very messy. You can end up with lots of blank lines and a couple of hard-to-spot assignment groups way down the bottom of the scroll area.

The solution is to add one extra line to the calculations section for the “operator” format control. A reverse-alphabetic sort means that the blank lines end up at the end.

```
$sort.junk=rtecall("sort", $sort.result, assignment.groups in $file, 0, 1)
```

14 Installing the Web tier on ServiceManager 9.2 on Linux

If you want error messages, you will need to create a file called `sm.log` in the top level directory of tomcat; HP doesn't document this anywhere as far as I can see.

On a RedHat system, using the tomcat 5.5 RPM, do this:

```
touch /var/log/tomcat5/sm.log
chown tomcat /var/log/tomcat5/sm.log
ln -s /var/log/tomcat5/sm.log /usr/share/tomcat5/sm.log
```

If you get an error message about being unable to initialise EasyDOM, and you find the `sm.log` file says that EasyDOM is trying to load `org.apache.xalan.processor.TransformerFactoryImpl` but can't, (and so throws `javax.xml.transform.TransformerFactoryConfigurationError`), then this might be fixable by disabling `xml-common-apis` in tomcat.

I think this is because xalan in tomcat5 is only needed if you are running in Java 1.4 - I think it got brought into Java after that.

```
cd /usr/share/tomcat5/common/endorsed
mv \[xml-commons-apis\].jar \[xml-commons-apis\].disabled
```

15 DVD tricks

A common scenario is that you want to have a field which is read-only if a flag "active" is false.

You can create a DVD which says:

```
Read-only    [active] = false
condition:
```

But that will make the search screen read-only as well. Do this instead:

```
Read-only    [$L.mode] <> "search" & [active] = false
condition:
```

16 What to think about when bringing a new team or Service Desk to a ServiceManager environment

This is a disorganised list of notes of things to think about or do if a team is going to start using DSDP. It is collated from feedback I've received in delivering training these are the I wish someone had told me this earlier stuff.

Create a list of incident handling groups What teams do you have at the moment? What do they do?

Create a list of knowledge article categories You currently have documents somewhere that help your support staff do their jobs. Do you have several different repositories? Or just one per application? Who is allowed to modify them? Who is allowed to read them?

Create a list of user profiles What different roles do your staff do? Each user can have only one user profile, so you might need several profiles. Things to consider in making these user profiles:

- *Incident role.* Can people with this user profile allocate tickets to other people, and pick things out of the queue for themselves to work on (either makes them an incident coordinator)? Or are all tickets given to them (incident analyst)?
- *Change roles.* Do they record changes? Do they review changes? Do they recommend changes?
- What incident handling group are they in?
- Can people with this user profile publish knowledge articles? Or can they only author them (so that someone else has to publish them)? Or do they have no rights to write anything?

Create a list of staff in your team Assign each a user profile. Create a spreadsheet of these or keep it in some convenient format which will be easy to load later.

Create a list of things which you support These will be *bizservice configuration items* – these can be chosen as the “Affected Service”. What is the escalation path for these things? Which teams will get involved and when?

Define the service level objectives

Create a list of the top 10 common incidents These are the ones that you want the service desk to be able to deal with quickly. These will often turn into templates. For each one of these, make sure there is a procedure for handling them. Usually knowledge articles are appropriate here.

Create knowledge articles Create them about anything and everything that you want the service desk (level 1s) to deal with. You might find this easier if you've been on training first. Create some more articles for your level 2 staff for all the procedures you want to follow.

Create your Service Catalogue items Come up with as many common changes and requests as you can. Each thing that you are asked for is going to have:

- Information that needs to be collected when the user calls or requests this thing.
- Approvals that need to be given
- Tasks that need performing (and who is going to do them).

Create your configuration item database You have a list of things (Software modules? Computers? Printers?) which you support. Do you need to be able to report on the incidents and changes done to these things? (If you don't, just ignore this stanza.)

- How are you going to initially import these lists of things?
- How are you going to keep it maintained? Are you receiving a feed from another application? Are you receiving a feed from several applications? If so, how are you going to reconcile differences?
- What attributes are you going to maintain?
- Do you have baselines that these things should conform to?
- Are you going to audit?
- Are there any business service configuration items?
- What are you going to do with these incident and change analysis reports?

17 Last-fortnight Training Requirements

The last two weeks before a Service roll-out are the most crucial and busy time for training staff. The bulk of what will be required for a is expressed in table 3 on page 20.

Some assumptions:

- Team leaders and managers (who might not necessarily use ServiceManager themselves for much other than approvals and reporting) will still want training along with their staff so that they know what is going on.

- Release control, configuration management, integrations, development and general system administration will be done by small teams. They would usually attend longer, more in-depth classes.

	Service Desk	Level 2 / 3 Techs	Knowledge management team	Regional support staff	Change management team
Creating dashboards and views. Navigation. Alt-F9 and filling. Approx 30min	✓	✓	✓	✓	✓
Logging and closing interactions. Using existing knowledge articles. Approx 1hr	✓	✗	✗	✗	✗
Understanding the incident workflow; updating and closing incidents. Searching and creating views. Using the mobility client for incidents. Approx 3hrs	✓	✓	✗	✓	✗
Writing knowledge articles, the knowledge document workflow Approx 1hr	✗	✓	✓	✗	✗
Change: normal, standard and service catalogue changes. Delegating and performing tasks. Using the mobility client for change. Approx 3hrs	✗	✓	✗	✗	✓
Problem management: creating, managing and closing problems; creating, managing and closing known errors. Approx 1hr	✗	✓	✗	✗	✗

Table 3: Standard Training Matrix