



smf(5): Solaris 10 Service Management Facility

Ambreesh Khanna
Chief Technologist, x64 Servers and Solaris
US Client Solutions
Sun Microsystems



smf (5) Service

- What's a service?
 - Abstract description of a long-lived software object
 - Each instance of a service has a well-defined state and a well-defined error boundary [process contract]
 - Each service defines “methods” and “dependencies”
 - Start, stop, refresh, etc.; interservice relationships
- A consistent specification
 - Can state dependency characteristics (grouping, restart)
 - Generic restart facility provided by default; customized restart capabilities available to vendor
- Admins can get a meaningful *system* view

Motivation

- Acknowledge difference between a service and a mere program
- No operating system support for service-based management
- No knowledge of service boundary and interservice relationships
- “Thousands of text files” not a design principle
- Take advantage of faster hardware
- Where can we remove opportunities for error from the system?

SMF Identifier

- FMRI – Fault Management Resource Identifier
 - `svc://localhost/network/login:rlogin`
 - Absolute path
 - `svc://network/login:rlogin`
 - Path relative to local machine
 - `network/login:rlogin`
 - `rlogin`
 - `svc://localhost/system/cryptosvc:default`
 - `svc://system/cryptosvc:default`
 - `lrc:/etc/rc3_d/S90samba`

SMF Identifier

- FMRI – Fault Management Resource Identifier
 - Functional categories
 - Application
 - Device
 - Milestone
 - Network
 - Platform
 - Site
 - System

Service States

- online
 - instance enabled and started
- offline
 - instance enabled but not yet running
- disabled
 - instance not enabled and not running
- maintenance
 - instance encountered an error, manual intervention required

Service States

- legacy_run
 - used by legacy services
 - service can be observed by not managed by SMF
- degraded
 - instance enabled but running at limited capacity
- uninitialized
 - initial state for all services before their configuration has been read

svcs (1) in action

- Pre-SMF “service” characteristics
 - what does sendmail depend on?
 - what depends on sendmail?
 - which processes constitute the mail service?
 -

```
$ pgrep -lf sendmail
708  /usr/lib/sendmail -bd -q15m
685  /usr/lib/sendmail -Ac -q15m

$ find /etc -name \*sendmail\* -print
/etc/rc0.d/K36sendmail
/etc/rc1.d/K36sendmail
/etc/rc2.d/S88sendmail
/etc/rcS.d/K36sendmail

$ less /etc/rc2.d/S88sendmail
....
```


svcs (1) in action

- List active instances, sorted by state, time
- Show dependencies (-d) and dependents (-D)
- Show member processes (-p), additional details (-v)

```

$ svcs
STATE          STIME      FMRI
....
online         18:18:30   svc:/network/http:apache
online         18:18:29   svc:/network/smtp:sendmail
....
$ svcs -p network/smtp:sendmail
STATE          STIME      FMRI
online         18:18:29   svc:/network/smtp:sendmail
               18:18:29   100180 sendmail
               18:18:29   100181 sendmail

$ svcs -v network/smtp:sendmail
STATE          NSTATE      STIME      CTID      FMRI
online         -           18:18:29   21        svc:/network/smtp:sendmail

$ svcs -d network/smtp:sendmail
STATE          STIME      FMRI
online         18:17:44   svc:/system/identity:domain
online         18:17:52   svc:/network/service:default
....

```

svcs (1) in action

- List active instances, sorted by state, time
- Show dependencies (-d) and dependents (-D)
- Show member processes (-p), additional details (-v)

```
$ svcs -D network/physical
STATE          STIME          FMRI
disabled       Nov_24         svc:/network/dns/client:default
disabled       Nov_24         svc:/network/dns/server:default
disabled       Nov_24         svc:/network/rarp:default
disabled       Nov_24         svc:/network/rpc/bootparams:default
disabled       Nov_24         svc:/network/slp:default
disabled       Nov_24         svc:/network/shell:kshell
online         Nov_24         svc:/application/print/cleanup:default
online         Nov_24         svc:/system/identity:node
online         Nov_24         svc:/system/identity:domain
online         Nov_24         svc:/network/initial:default
online         Nov_24         svc:/milestone/single-user:default
online         Nov_24         svc:/network/inetd:default
online         Nov_24         svc:/network/nfs/client:default
online         Nov_24         svc:/network/shell:tcp
online         Nov_24         svc:/network/shell:tcp6only
online         Nov_24         svc:/network/nfs/server:default
```

```
$
```

svcs (1) in action

- Diagnose instances in unusual state (-x)

```
$ svcs -x
svc:/application/print/server:default (LP Print Service)
  State: disabled since Thu 30 Sep 2004 01:14:16 PM PDT
  Reason: Disabled by an administrator.
    See: http://sun.com/msg/SMF-8000-05
    See: lpsched(1M)
  Impact: 1 service is not running.

svc:/system/metainit:default (Solaris Volume Manager(SVM) Init service.)
  State: maintenance since Thu 30 Sep 2004 01:14:16 PM PDT
  Reason: Completes a dependency cycle.
    See: http://sun.com/msg/SMF-8000-HP
    See: metainit(1M)
  Impact: 0 services are not running.
```

svcs (1) in action

- List service details (-l)

```
$ svcs -l network/smtp:sendmail
fmri          svc:/network/smtp:sendmail
enabled      true
state        online
next_state   none
restarter    svc:/system/svc/restarter:default
contract_id  46
dependency   require_all/refresh file://localhost/etc/nsswitch.conf (-)
dependency   require_all/refresh file://localhost/etc/mail/sendmail.cf (-)
dependency   optional_all/none svc:/system/system-log (online)
dependency   require_all/refresh svc:/system/identity:domain (online)
dependency   require_all/refresh svc:/milestone/name-services (online)
dependency   require_all/none svc:/network/service (online)
dependency   require_all/none svc:/system/filesystem/local (online)

$ svcs -l system/metainit
fmri          svc:/system/metainit:default
name         Solaris Volume Manager(SVM) Init service.
enabled      false
state        maintenance
next_state   none
restarter    svc:/system/svc/restarter:default
dependency   require_all/none svc:/system/filesystem/minimal (online)
dependency   require_all/none svc:/system/identity:node (online)
```

svcadm (1M) in action

- pre-SMF disabling of a “service”

```
$ pgrep -lf sendmail
708 /usr/lib/sendmail -bd -q15m
685 /usr/lib/sendmail -Ac -q15m

$ pkill sendmail

.....reboot.....

$ pgrep -lf sendmail
355 /usr/lib/sendmail -bd -q15m
276 /usr/lib/sendmail -Ac -q15m
```

- No log of why sendmail went down
- Reboot will restart sendmail, is this desired?

```
$ mv /etc/rc2.d/S88sendmail /etc/x/S88sendmail

.....apply patch/upgrade.....

$ ls /etc/rc2.d/S88sendmail
/etc/rc2.d/S88sendmail
```

svcadm (1M) in action

- Enable, disable, refresh, restart service instances
- Mark in special states (maintenance, degraded)

```
$ grep ambreesh /etc/user_attr
ambreesh:::auths=solaris.smf.modify,solaris.smf.manage
```

```
$ svcs -a network/http:apache
STATE          STIME      FMRI
uninitialized  19:17:33  svc:/network/http:apache
```

```
$ svcadm enable network/http:apache
STATE          STIME      FMRI
online         19:19:01  svc:/network/http:apache
```

```
$ vi /etc/apache/httpd.conf
```

```
$ svcadm refresh network/http:apache
```

```
$ svcs -a network/http:apache
STATE          STIME      FMRI
online         19:19:33  svc:/network/http:apache
```

```
$ svcadm disable network/http:apache
```

```
$ svcs -a network/http:apache
STATE          STIME      FMRI
disabled       19:20:07  svc:/network/http:apache
```

Service Manifest

- Each package delivering services does so via a “service manifest”
- xml file containing description of service
 - `/var/svc/manifest`
 - Dependencies on other services and methods for service instance start/stop/refresh
 - Default properties and “service template”, which provides support for administrative apps via
 - Localized property descriptions
 - Links to documentation
 - Soon: meaningful property values (valid ranges, definitions, etc.)
 - See `/usr/share/lib/xml/dtd/service_bundle.dtd.1`

coreadm (1M) service description

```
<service name='system/coreadm' type='service' version='1'>
  <create_default_instance enabled='false' />
  <single_instance />
  <dependency name='usr' grouping='require_all'
    restart_on='none' type='service'>
    <service_fmri value='svc:/system/filesystem/minimal' />
  </dependency>
  <exec_method type='method' name='start'
    exec='/usr/bin/coreadm -u' timeout_seconds='3' />
  <exec_method type='method' name='stop'
    exec=':true' timeout_seconds='0' />
  <property_group name='startd' type='framework'>
    <propval name='duration' type='astring' value='transient' />
  </property_group>

  <stability value='Unstable' />
  <template>
    <common_name><loctext xml:lang='C'>
      System-wide core file configuration service.
    </loctext></common_name>

    <documentation>
      <manpage title='coreadm' section='1M'
        manpath='/usr/share/man' />
    </documentation>
  </template>
</service>
```


SMF repository

- All manifests stored in persistent, transaction-based repository
 - Transactions/snapshots allow “undo”, rollback to safe configuration
 - Repository can be local, in directory [later], or mixed [later]
- **NOT** a giant registry: mainly svc mgmt properties
- Data imported either at boot time from manifests or via `svccfg import`
- `svccfg` (1M) used to manipulate repository data
- `svcprop` (1M) used to view repository data
- `svc.configd` (1M) is the repository daemon

svccfg (1M) in action

- View properties of a service

```
$ svccfg
svc:> list
....
network/rpc/nisplus
....
platform/i86pc/kdmconfig
network/fs/tcp6
svc:> select network/ssh
svc:/network/ssh> select default
svc:/network/ssh:default> listprop
general                               framework
general/package                       astring  SUNWsshdr
general/enabled                       boolean  true
restarter                             framework  NONPERSISTENT
restarter/contract                    count    280
restarter/start_pid                   count    7756
restarter/auxiliary_state             astring  none
restarter/next_state                  astring  none
restarter/state                       astring  online
restarter/state_timestamp              time     1101530796.234165000
```

svccfg (1M) in action

- Add a new property “myprop”

```

svc:/network/ssh:default> setprop general/myprop=astring:"demoval"
svc:/network/ssh:default> listprop
general                framework
general/package       astring  SUNWsshdr
general/enabled       boolean  true
general/myprop        astring  demoval
restarter             framework  NONPERSISTENT
restarter/contract    count    280
restarter/start_pid   count    7756
restarter/auxiliary_state astring  none
restarter/next_state  astring  none
restarter/state       astring  online
restarter/state_timestamp time     1101530796.234165000
svc:/network/ssh:default> delprop general/myprop
svc:/network/ssh:default> listprop
general                framework
general/package       astring  SUNWsshdr
general/enabled       boolean  true
restarter             framework  NONPERSISTENT
restarter/contract    count    280
restarter/start_pid   count    7756
restarter/auxiliary_state astring  none
restarter/next_state  astring  none
restarter/state       astring  online
restarter/state_timestamp time     1101530796.234165000

```

SMF Components

- Master Restarter
 - `svc.startd(1M)`
 - responsible for starting and restarting services
 - starts services when their dependencies are met
 - restarts failed services
 - shuts down services when dependencies no longer met

```
$ pgrep -lf ssh
7759 /usr/lib/ssh/sshd

$ pkill ssh

$ pgrep -lf ssh
8635 /usr/lib/ssh/sshd
```

SMF Components

- Delegated Restarter
 - restarter for a related set of services

SMF Components

- Delegated Restarter
 - `inetd` (1M) the only one so far
 - responsible for internet services
 - `/etc/inet/inetd.conf` now deprecated
 - `inetadm` (1M) manages internet services
 - `inetconv` (1M) converts `inetd.conf` into `inetd`

```

$ svcs -R network/inetd:default
STATE          STIME      FMRI
disabled      Sep_30    svc:/network/rpc/meta:tcp
...
online        Oct_01    svc:/network/rpc/ttdbserver:tcp6
online        Oct_03    svc:/network/shell:tcp
online        16:58:06  svc:/network/finger:default
online        19:11:06  svc:/network/rpc/rstat:udp
online        23:08:21  svc:/network/rpc/rusers:udp
offline       Sep_30    svc:/application/print/rfc1179:default
$ svcs -R network/inetd:default | wc -l
66

```

inetadm (1M) in action

- List internet services and view/modify their properties

```

$ inetadm
ENABLED STATE FMRI
enabled online svc:/network/rpc/gss:ticotsord
disabled disabled svc:/network/tname:default
enabled online svc:/network/security/ktkt_warn:ticotsord
enabled online svc:/network/telnet:default
.....
disabled disabled svc:/network/apocd/udp:default
disabled disabled svc:/network/uucp:default
disabled disabled svc:/network/security/krb5_prop:tcp
enabled online svc:/network/rpc-100235_1/rpc_ticotsord:ticotsord
enabled online svc:/network/rpc-100424_1/rpc_ticotsord:ticotsord
enabled online svc:/network/rpc-100083_1/rpc_tcp:tcp
enabled online svc:/network/rpc-100083_1/rpc_tcp:tcp6
enabled online svc:/network/rpc-100068_2-5/rpc_udp:udp
enabled online svc:/network/rpc-100068_2-5/rpc_udp:udp6
enabled online svc:/network/fs/tcp6:default

$inetadm -e network/tname

$inetadm | grep tname
enabled online svc:/network/tname:default

```

inetadm (1M) in action

- List internet services and view/modify their properties

```
$ inetadm -l network/telnet
SCOPE      NAME=VALUE
           name="telnet"
           endpoint_type="stream"
           proto="tcp6"
           isrpc=FALSE
           wait=FALSE
           exec="/usr/sbin/in.telnetd"
           user="root"
default   bind_addr=""
default   bind_fail_max=-1
default   bind_fail_interval=-1
default   max_con_rate=-1
default   max_copies=-1
default   con_rate_offline=-1
default   failrate_cnt=40
default   failrate_interval=60
default   inherit_env=TRUE
default   tcp_trace=FALSE
default   tcp_wrappers=FALSE
```


inetadm(1M) in action

- List internet services and view/modify their properties

```
$ inetadm -m network/telnet tcp_trace=TRUE
```

```
$ inetadm -l network/telnet
```

```
SCOPE      NAME=VALUE  
           name="telnet"  
           endpoint_type="stream"  
           proto="tcp6"  
           isrpc=FALSE  
           wait=FALSE  
           exec="/usr/sbin/in.telnetd"  
           user="root"  
default   bind_addr=""  
default   bind_fail_max=-1  
default   bind_fail_interval=-1  
default   max_con_rate=-1  
default   max_copies=-1  
default   con_rate_offline=-1  
default   failrate_cnt=40  
default   failrate_interval=60  
default   inherit_env=TRUE  
           tcp_trace=TRUE  
default   tcp_wrappers=FALSE
```

Troubleshooting

- Corrupted repository

```
$ svc.configd: smf(5) database integrity check of:
```

```
  /etc/svc/repository.db
```

```
failed. The database might be damaged or a media error might have prevented it from being verified. Additional information useful to your service provider is in:
```

```
  /etc/svc/volatile/db_errors
```

```
The system will not be able to boot until you have restored a working database. svc.startd(1M) will provide a sulogin(1M) prompt for recovery purposes. The command:
```

```
  /lib/svc/bin/restore_repository
```

```
can be run to restore a backup version of your repository. See http://sun.com/msg/SMF-8000-MY for more information.
```

Troubleshooting

- Restore procedure

```
root@vitalstatistix# /lib/svc/bin/restore_repository
```

```
Repository Restore utility
```

```
See http://sun.com/msg/SMF-8000-MY for more information on the use of  
this script to restore backup copies of the smf(5) repository.
```

```
If there are any problems which need human intervention, this script will  
give instructions and then exit back to your shell.
```

```
Note that upon full completion of this script, the system will be rebooted  
using reboot(1M), which will interrupt any active services.
```

```
The following backups of /etc/svc/repository.db exist, from  
oldest to newest:
```

```
manifest_import-20050204_142224  
manifest_import-20050204_173653  
manifest_import-20050206_175253
```

```
.....
```

Troubleshooting

- Restore procedure

```
manifest_import-20050209_122006  
boot-20050313_201600  
boot-20050314_095635  
boot-20050314_174541  
boot-20050315_112629
```

The backups are named based on their type and the time what they were taken. Backups beginning with "boot" are made before the first change is made to the repository after system boot. Backups beginning with "manifest_import" are made after `svc:/system/manifest-import:default` finishes its processing. The time of backup is given in `YYYYMMDD_HHMMSS` format.

Please enter one of:

- 1) boot, for the most recent post-boot backup
- 2) manifest_import, for the most recent manifest_import backup.
- 3) a specific backup repository from the above list
- 4) -seed-, the initial starting repository. (All customizations will be lost.)
- 5) -quit-, to cancel.

```
Enter response [boot]: -quit-  
Exiting.  
root@vitalstatistix:/etc/svc#
```

Boot process

- milestones analogous to run-levels
 - S milestone/single-user:default
 - milestone/name-service
 - 2 milestone/multi-user:default
 - 3 milestone/multi-user-server:default
- Boot without any services enabled, then enable all services

```
ok boot -m milestone=none
```

```
login as root
```

```
# svcadm milestone -t all
```

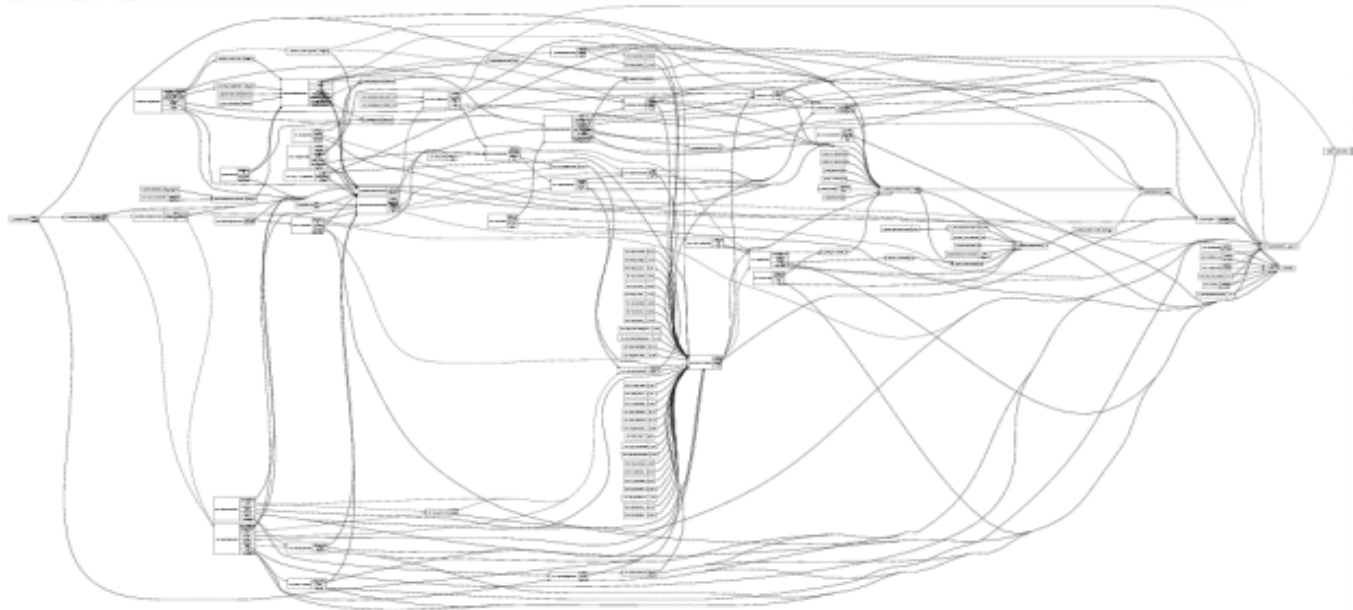
Boot process

- Verbose boot available at boot prompt
 - Persistently verbose? Set options/logging to verbose on system/svc/restarter:default

```
Select (b)oot or (i)nterpreter: b -mverbose
SunOS Release 5.10 Version smf-mdb-on10 32-bit
Copyright 1983-2004 Sun Microsystems, Inc. All rights reserved.
Use is subject to license terms.
DEBUG enabled
WARNING: consconfig: could not find driver for screen device /isa/display@1,3b0
[ network/loopback:default starting (Loopback network interface) ]
[ network/pfil:default starting (pfil) ]
[ system/filesystem/root:default starting (Root filesystem mount) ]
[ network/physical:default starting (Physical network interfaces) ]
Oct  6 12:21:15/14: system start time was Wed Oct  6 12:21:10 2004
[ system/filesystem/usr:default starting (/usr and / mounted read/write) ]
[ system/identity:node starting (system identity (nodename)) ]
Hostname: trolls-b10
[ system/device/local:default starting (Standard Solaris device configuration.) ]
[ system/filesystem/minimal:default starting (Local filesystem mounts) ]
[ milestone/devices:default starting (Device configuration milestone.) ]
[ system/identity:domain starting (system identity (domainname)) ]
[ system/cryptosvc:default starting (Cryptographic services) ]
[ system/manifest-import:default starting (Service manifest import) ]
[ system/sysevent:default starting (System event notification service.) ]
NIS domain name is mpklab.sfbay.sun.com
[ system/coreadm:default starting (System-wide core file configuration service) ]
....
```

Boot process

- `init` reads `/etc/inittab`
 - `smf::sysinit:/lib/svc/bin/svc.startd >/dev/msglog 2<>/dev/msglog </dev/console`
- starts up master restarter





Ambreesh Khanna
ambreesh.khanna@sun.com

Developer Day

