

Inside Solaris™

Tips & Techniques for users of Sun Solaris

Using TCM for modeling and designing software

by H-W Schlote

Modern tasks in the computer industry are complex, and solving them most often requires a professional approach to creating design papers, producing specifications, and writing documentation. The days are (or at least should be) long over when it was sufficient to present a pro-

gram (the binary or executable) and a couple of source code lines to solve a task.

Tools are needed for creating pictures and graphs—for structured approaches as well as for object-oriented software design only. There are commercial tools like Rational Rose. Tools like these have the disadvantage of being incredibly expensive. Additionally, Rational Rose can handle object-oriented software only. There are a few other commercial tools for conceptual modeling, but most of them can be used on MS-based operating systems only.

In this article, we'll introduce TCM, the toolkit for conceptual modeling, which is free software—at least for non-commercial use. On Solaris, TCM is very stable. Unfortunately, on HP-UX core dumps were encountered, but even versions of Rational Rose show the same behavior on HP-UX.

What's TCM?

TCM was developed by the faculty of Mathematics and Computer Science of the Free University of Amsterdam (Vrije Universiteit Amsterdam). TCM isn't just one tool but a set of graphical editors for a number of graphical notation systems used in software specification methods. See **Figure A** for the main menu.

TCM's tools

Besides having generic diagram, table, and tree editors, TCM also consists of tools for:

- Entity Relationship Diagrams
- Class Relationship Diagrams

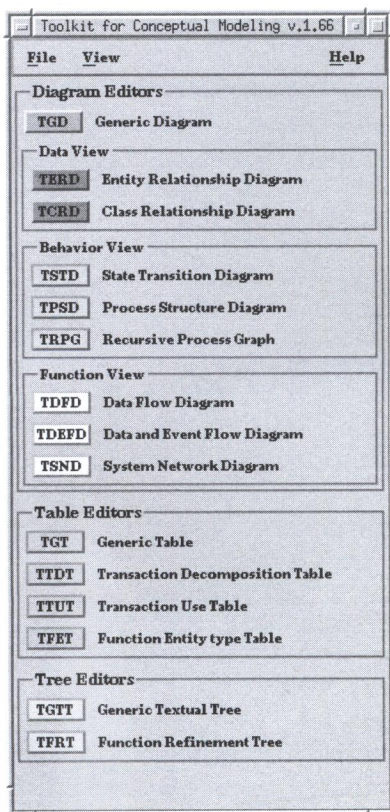


Figure A: This is the TCM main menu.

In this issue:

1 Using TCM for modeling and designing software

4 Secure networking with SSL

7 Restoring your Sun

10 **Net.Update:** Organizations that make the Internet work

12 Managing environment variables with Env

16 **Solaris Q & A:**
• Resizing partitions

• Fixing network settings

- State Transition Diagrams
- Recursive Process Graphs (also called life cycle diagrams)
- Process Structure Diagrams
- Data Flow Diagrams
- Data and Event Flow Diagrams
- System Network Diagrams
- Transaction Decomposition Tables
- Transaction-Use Tables
- Function-Entity Type Tables
- Function Refinement Trees

There are 12 tools plus 3 generic ones. Actually the generic tools are supersets of the specific ones with regard to the graphical objects available. You could argue that the specific tools aren't necessary because you can draw a diagram created with a specific tool with the corresponding generic one, too. But the specific tools support specific constraint checking.

For example, in the Entity Relationship Tool an entity type can be connected with a relation through a functional relation only. Thus it's ensured that the correct diagrams will be created. We'll cover constraint checking with TCM in detail later in this article.

The Entity Relationship Diagram

Figure B shows a simple Entity Relationship Diagram for planes constructed with the specific TCM tool. On the left side of the plane entity, the three possible types of aircraft are joined by a disjunctive taxonomic junction, hence the letter *d* inside the junction. In air guidance systems, planes are classified into categories heavy, medium, and light. Thus, the taxonomic junction is classified *de* to show its disjunctive and exhaustive nature. Allowed classifications for taxonomic junctions in Entity Relationship Diagrams are *d*, *e*, and *de*. For example, a B747 on a scheduled flight from Frankfurt to New York would match entities scheduled and heavy.

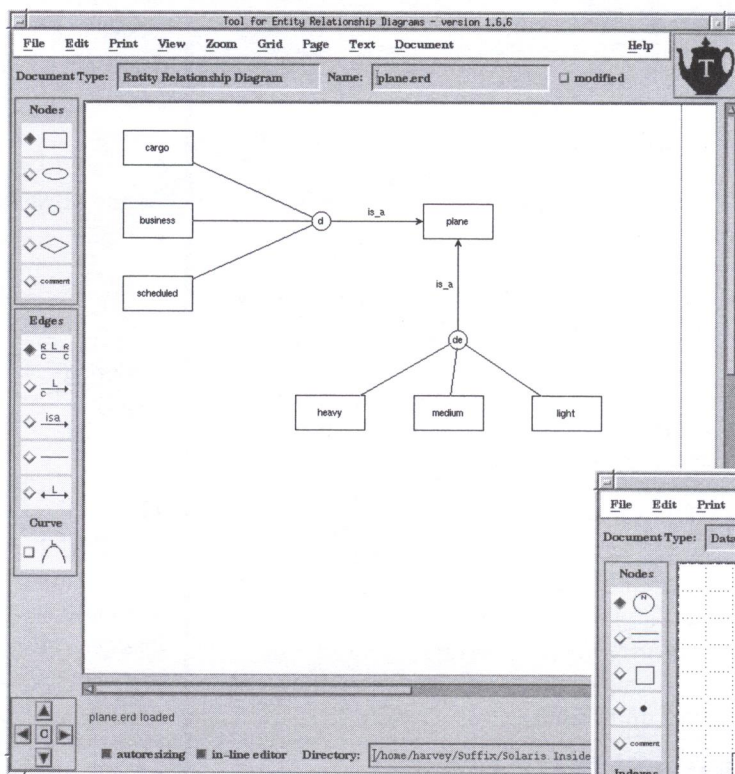


Figure B: Here you can see a sample of an Entity Relationship Diagram built in TCM.

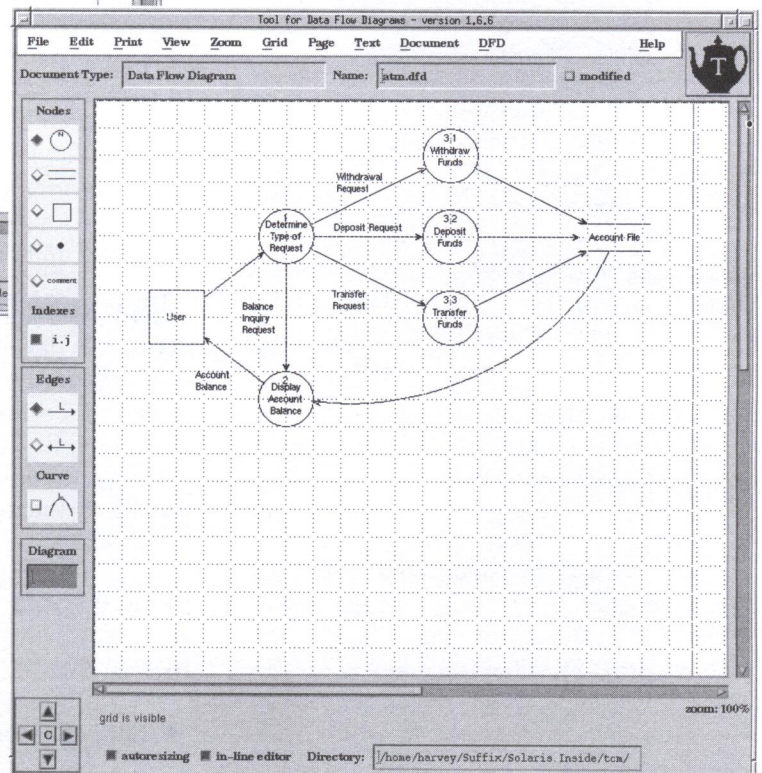


Figure C: The DFD for our ATM application is shown here.

The in-line editor and Motif-style text box

TCM supports two kinds of editors: an in-line editor and a Motif-style text box. For initially creating a diagram, the in-line editor is our first choice, because naming nodes and edges can be performed faster with it. For changing a diagram, the Motif-style text box is very useful. But when to use which editor is more a question of personal taste.

Straight, segmented, and curved edges

By default, TCM draws straight edges between nodes. Sophisticated algorithms are implemented into TCM for distributing multiple straight edges connecting the same pair of nodes equally. In addition to straight edges, TCM supports segmented and curved edges. A segmented edge is drawn if intersection points are set with the middle mouse button while connecting two nodes. To create an edge drawn as a Bezier curve (curved edge), you merely need to toggle the corresponding button labeled Curve.

The edge connecting data store Account File with data process Display Account Balance, in the dataflow diagram shown in Figure C, is a Bezier curve. In this diagram, data flow for an ATM application is shown. Note the sub-indexing in nodes Withdraw Funds, Deposit Funds, and Transfer Funds.

Three constraints

TCM supports three kinds of constraints. The first, *built-in constraints* are constraints that can never be violated because there's no command in the user interface to achieve that. For example, you can't put a class object into an Entity Relationship Diagram.

The second type is *immediately enforced constraints*. When you perform a command that would violate a constraint that's immediately enforced, this command is rejected by TCM and a pop-up window with an error message displays. This happens, for example, if the user tries to connect a relation with an entity through an empty edge.

And the final type, *soft constraints*, are constraints that can be violated. Soft constraints are checked by TCM when the Check Document from the Document menu is issued. Check Document displays a list of error messages in a pop-up window. As opposed to the previous two classes of constraints, the user is

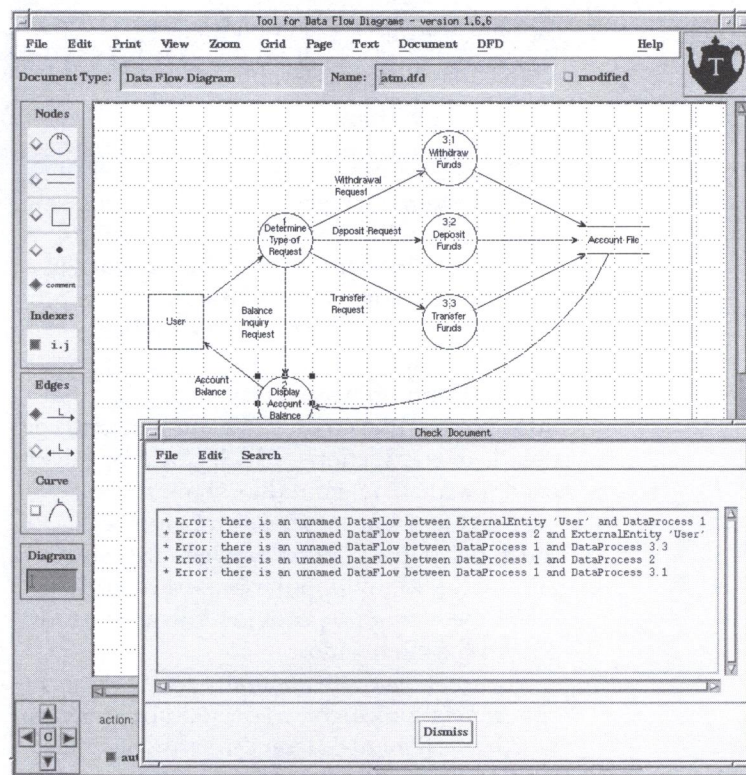



Figure D: This is the Data Flow Diagram showing the soft constraint violation message window.

responsible for correcting the diagram. An example of a violated soft constraint is an unnamed edge. Figure D shows the Data Flow Diagram with the message display window resulting from the Check Document action.

Conclusion

TCM runs on most UNIX systems with X Windows. Thus, there's no necessity to switch to a Windows desktop when you want to create a graph or diagram. Instead, you can create diagrams—for example, for documentation purposes—accompanying your development process on the same machine where you're performing your daily tasks. You can do so with the stability of an operating system you're used to. 

References

- TCM's Web site: www.wis.cs.utwente.nl:8080/~tcm/index.html
- *User's guide and reference for TCM*, F. Dehne and R. J. Wieringa
- *Object-Oriented Software Development*, M. Lorenz, Prentice Hall, 1993

Secure networking with SSL

by Paul A. Watters

Solaris is often selected as the operating system of choice for implementing secure, reliable eCommerce solutions. A central requirement for Web-based transactional processing is the assurance of data protection and reliability for online exchanges of sensitive data.

One common way to protect confidential information, like credit card numbers and/or details of transactions, which could be misused for financial gain (like share transactions), is to use a secure-socket layer (SSL) in the network connection between the client and server. Many popular browsers, such as Netscape Navigator, provide full support for SSL on the client side.

Fortunately, there are many freeware and commercial products supported under Solaris which implement SSL on the server side. In this article, we'll examine the implementation of a secure application server solution using Apache, the most popular Web server, and a free version of SSL (SSLey, which has recently been released as OpenSSL). Apache and SSLey are interfaced using `mod_ssl`, an Apache module that connects the Web server to the SSL library.

Secure networking

In the media, we're constantly bombarded with accounts of electronic fraud using credit card numbers obtained through the Internet. While this kind of fraud no doubt occurs, it's surely easier for a potential thief to obtain credit card numbers by other means (that is, obtaining receipts for goods purchased by someone else's credit card). Given the substantial technological prerequisites for packet sniffing, and the kind of heuristic algorithms required to extract plain-text credit card from the gigabytes of data transferred around local area networks, it seems a less than profitable enterprise.

However, there may well be certain kinds of data that criminals would be prepared to expend large amounts of time and computing resources to extract from plain-text network transmissions. For example, a stockbroking company may circulate internal buy/sell orders for securities on the LAN between brokers, using a Web-based interface between

each broker's client and a central server. If this kind of data were illegally obtained by a competitor, even a few minutes advance notice of a certain kind or quantity of orders placed by a client could be misused.

Alternatively, passwords transmitted through the network in clear-text for administrative access to the server could be sniffed, with the buy/sell orders manipulated to promote a financial loss for the company. Again, only a few minutes of access would be required to inflict serious damage.

These scenarios aren't science fiction; they represent real risks that networked enterprises face daily. Fortunately, there are a number of tools available for eCommerce solutions built around Solaris, which can reduce (but not totally eliminate) the risk of information being obtained illegally. Many of these tools involve some kind of encryption technology, in either the exchange or encoding of sensitive data.

For example, the popular public-key cryptography software known as PGP (Pretty Good Privacy), created by Phil Zimmerman, encodes data by using a recipient's public key and a sender's private key to ensure that only the sender and recipient can retrieve the encoded data. It isn't known whether anyone has succeeded in cracking PGP, but it has been widely adopted in email clients.

Secure socket layer

Fortunately, for client-server computing adapted to the Web, there are also solutions available to dynamically encrypt data exchanged through the Web. The standard adopted by most organizations is the secure socket layer (SSL) protocol, which is currently in version 3.0. SSL aims to reduce the risk of data obtained or altered through packet sniffing and spoofing by using a two-tiered system of transport.

First, the SSL Record Protocol sits on top of existing Internet data exchange layers (for example, TCP), and encapsulates higher-level protocols. Second, the SSL Handshake Protocol facilitates authentication by reaching an agreement between client and server on an encryption algorithm and other parameters prior to the exchange of application data. SSL has a distinct advantage over PGP in that clients

and servers don't need to have exchanged public keys prior to a connection being made. The network layering involved in SSL is shown in **Figure A**.

The encryption algorithm used by both client and server depends on individual circumstances (for example, the highly publicized restrictions on exporting encryption algorithms from various nations). Fortunately, there are many algorithms and implementations that are available internationally to support the international focus of eCommerce on the Internet.

In this article, we present a solution for serving applications through the Internet with an add-in module for the freely-available Apache Web server, known as `mod_ssl`, and a free implementation of SSL, known as `SSLey` (now `OpenSSL`).

Installation and configuration

The installation and configuration of SSL support for Apache is fairly straightforward for Solaris, although a number of compilation options are available to meet individual requirements. We're compiling Apache 1.3.4 with `SSLey-0.9.0b` and `mod_ssl-2.1.7-1.3.4`. This isn't the latest version of all three packages, but it's a stable and tested combination for Solaris. More operational testing is required before shifting from `SSLey` to the newer `OpenSSL` implementation.

The important pre-installation stage involves identifying which version of each software package is compatible with the other. Like many other implementations of software associated with the WWW, `mod_ssl`, in particular, appears to change fortnightly or monthly. Fortunately, versions of `mod_ssl` are identified by two version numbers; one for the module revision (2.3.3) and one for the target Apache version (1.3.6). Thus, `mod_ssl-2.3.3-1.3.6` is the most current version of `mod_ssl` that's compatible with Apache 1.3.6.

You can obtain current versions of each software package from these distribution sites:

- www.apache.org (Apache Web server)
- www.openssl.org (OpenSSL)
- www.modssl.org (`mod_ssl`)

The install file that accompanies the `mod_ssl` tarball contains a concise overview of the

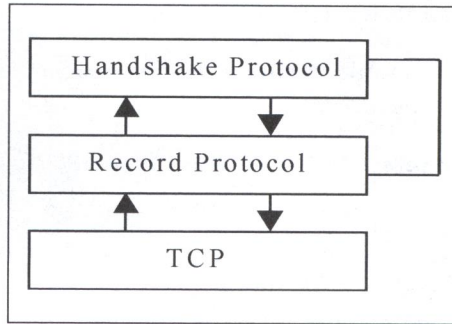


Figure A: The SSL protocol consists of a Handshake and Record Protocol that sits on top of Internet protocols like TCP.

steps required to install and configure the sources. To begin installing and configuring the sources, first unpack the sources for Apache, `mod_ssl`, and `SSLey` into a source directory (`/usr/local/src` for example). Now, ensure that the Apache and `SSLey` distributions are visible in the parent directory for `mod_ssl`. Depending on your location, you may also need to install the `RSAREf` library.

Next, change into the source directory for `mod_ssl`, and enter the command:

```
./configure\
--with-apache=../apache_1.3.4 \
--with-ssley=../SSLey-0.9.b \
--prefix=/usr/local/pkgs/apache
```

Of course, if your versions of `mod_ssl` and `SSLey` are different from those shown, the appropriate source directories in `mod_ssl`'s parent directory should be substituted. The configuration script produces output explaining the changes made to the Apache configuration file:

```
Configuring mod_ssl/2.1.7 for Apache/1.3.4
+ Apache location: ../apache_1.3.4
  ➔(Version 1.3.4)
+ Auxiliary patch tool: ./etc/patch/
  ➔patch (local)
+ Applying packages to Apache source tree:
[SSL Module Source]
[SSL Support]
[SSL Configuration Additions]
[SSL Module Documentation]
Done: mod_ssl source extension and
  ➔patches successfully applied.
```

Next, change into the Apache source directory:

```
$ cd ../apache_1.3.4
```

The usual compilation messages for Apache appear on the screen:

```
bash-2.01# make
==> src
==> src/os/unix
==> src/ap
==> src/main
==> src/modules
```

including the compilation of the SSL module:

```
==> src/modules/ssl
```

All packages and modules should compile successfully under Solaris 7. Next, a test configuration (including an X.509 certificate) can be generated by issuing the command:

```
$ make certificate
```

In the samples we provide, a certificate is generated for the company Snake Oil, CA, and an RSA private key (1024 bit) is also created. Prospective users should note the caveat generated in the output display:

```
WARNING: Do not use this for real-life/
production systems.
```

No doubt, there are many implementations of SSL and Apache online which have certificates for Snake Oil CA.

The SSL-aware Apache Web server can now be started with the following command:

```
$APACHE_HOME/bin/apachectl startssl
```


The default port 443 should now be listening for secure connections from clients. Of course, the configuration of the Apache Web server is a fairly detailed topic in itself. Refer to the Apache documentation for details.

Conclusion

The freely available combination of the Apache Web server, and a module that supports SSL, provides a powerful foundation for ensuring that transmitted network data is safer. Clearly, there are a number of constraints that affect the quality of performance of this system, including a trade-off between network bandwidth and the level of encryption desired (for example, 40 bit versus 128 bit and above).

Financial institutions typically make an investment in high-speed LANs and fast CPUs to compensate for the extra load of encryption and decryption, whereas most users feel safer with 40-bit encryption of credit card data as opposed to clear-text transmission.

It's important to note that OpenSSL isn't the only version of SSL that has been implemented with eCommerce or client-server computing in mind, and not all have been written in C. For example, there's a pure Java implementation called EspreSSL, which can be downloaded from www.vonneida.org/EspreSSL.

An interesting feature of this free software is that it's independent of any version or type of encryption technology. Instead, a plug-in facility has been developed so developers can include their own code for encryption, or use a third-party commercial version from RSA, or a patent-free version like Diffie-Hellman. 

ZD Journals
On the
WEB

Looking for more tips and techniques? There's plenty in store at www.zdjournals.com. Sign up for a FREE issue of another journal, check out our CD-ROM products, catch "The Daily Buzz," or sample a new article from one of our other industry leading journals. Plus, our online community gives you an opportunity to get in touch with our experts, as well as your fellow readers. Visit us today!

Restoring your Sun

DOWNLOAD

ftp.zdjournal.com/sun

by Werner Klauser

How often have you read a report or article on how to back up your system(s)? But do they ever say how the backups are to be used to restore your system? What's the purpose of using a simple cron job to steer your local tape drive or complex backup software using an FDDI connected tape silo if you need a UNIX software guru to restore your replaced system disk? Of course this is the disk that dies at seven in the morning, two hours before the year's most important business meeting. Then your boss has to inform the CEO why the system administrator has such a large salary.

This article will explain how to restore the operating system of a simple Solaris 2.x system using ufsdump backup tapes made on a locally connected tape drive. More complex scenarios are just variations of the basics.

Needed information

Restoring a system is only trivial if the backup is properly undertaken. And a backup isn't just the writing of data on tapes, it also means gathering some system information on paper to help when you need to restore it.

Which drive

First, you need to know which drive is the system or boot drive. For this article, we'll consider `/dev/dsk/c0t3d0s0` or simply just `c0t3d0s0` as being the system disk.

Disk slices or partitions

Next you need to know how the disk is sliced or partitioned. A good way to find out is to use the `format` command to print the partition table. As user root type:

```
# format
```

A list of the disks is shown and the user is asked to enter the system disk's number. Then the format menu is displayed:

```
format> partition
```

Next, the partition menu is displayed:

```
partition> print
```

Now disk partitioning information will be displayed. Copy the following information to a file for printing:

```
partition>quit
format> quit
#
```

This information, as well as the contents of `/etc/vfstab`, should be kept somewhere where it's easily accessible in case of a system failure.

The actual backup: what needs to be backed up

All the local file systems need to be backed up. This doesn't include `/proc`, `swap`, nor the disk drive `fd`. It's favorable to first back up the `/` (root) file system if multiple file systems are placed on a single backup medium.

ufsdump

`/usr/sbin/ufsdump -0uf /dev/rmt0mn` is used to backup the individual file systems. Dump level 0 is used so that the entire file system is dumped. The parameter `u` results in the update of the dump record stored in `/etc/dumpdates`. However, this information is never used. `f` followed by the dump file `/dev/rmt0mn` results in `ufsdump`'s data being written to the non-rewinding, medium-sized tape drive. The final parameter is the name of the file system's raw disk device name.

Backup script

The script `backup.sh` shown in Listing A on page 8, determines which file systems need to be backed up and writes these onto a non-rewinding tape drive. The cron job needs only to call this script to backup the system's file systems using `ufsdump`.

Restoring the operating system

The first order of business is to boot from the Solaris CD-ROM into standalone mode. Place the CD-ROM that was originally used to install the operating system into the system's CD-ROM drive. From the boot prompt (the OK prompt), type:


```
OK boot cdrom -s
```

This will boot a standalone shell directly off the CD-ROM.

Repartition the new disk

The disk needs to be partitioned using the `format` command. The name of the system or boot disk without the `/dev/dsk/` prefix is the `format` command's only parameter:

```
# format c0t3d0s0
```

Then the `format` menu is displayed. It isn't necessary to actually format the disk, since almost all SCSI disks are preformatted at the factory. Only the disk needs to be repartitioned. Select the partition option:

```
format> partition
```

Next, the `format` menu displays. Enter the following commands using the printout of the previous partition layout:

```
partition> 0
Enter partition id tag [root]: [Enter]
Enter partition permission flags
->[wm]: [Enter]
Enter new starting cyl [0]: [Enter]
Enter partition size [??b, ??c, ??mb, ??gb]:
-><number of cylinders>
partition>
```

Repeat this process for all the other slices on the disk. Be sure to zero out any unused slices on the disk to avoid problems later. When finished, label it with the following:

```
partition> label
Ready to label disk, continue? Yes
```

Now quit `format`.

```
partition> quit
format> quit
#
```

Re-create file systems

The next step is to create a new file system on the partitions that were created in the previous steps. A file system isn't needed on the swap partition:

```
# newfs /dev/rdisk/c0t3d0s0
```

Note that the raw disk `/dev/rdisk` is used for this step. Repeat this step for any other slices in the drive except slices 1 and 2. Slice 1 is usually used as swap space and doesn't need to be restored, nor does it require a file system. Slice 2 refers to the whole disk.

Restore the data

Next, you must mount the root partition so that it can be restored:

```
# mount /dev/dsk/c0t3d0s0 /a
```

Change the directory to `/a`:

```
# cd /a
```

Begin the restore process using `ufsrestore`. `r` stands for recursive, which restores the entire contents of the media into the current directory:

```
# ufsrestore rvf /dev/rmt/0mn
```

A lot of messages from `ufsrestore` will appear, followed by a list of files that are being extracted from the tape. Once the root file system has been restored, continue restoring any other file systems found on the replaced system disk. For example, suppose that `/var` is a separate file system on `c0t3d0s4` and that `/var` is the next `ufsdump` file on the tape. Then the following needs to be done:

```
# mount /dev/dsk/c0t3d0s4 /a/var
# cd /a/var
# ufsrestore rvf /dev/rmt/0mn
```

When all disk slices have been restored onto the replaced disk, eject the no-longer-used tape and then unmount the partitions, making sure to start at the leaves of the hierarchy tree:

```
# mt -t /dev/rmt/0m offline
# cd /
# umount /a/var
# umount /
```

Install boot block

After all the files have been restored to the replaced disk, it's necessary to install the boot block before the new disk can be used to boot the system. In order for the system to boot, the `boot(1M)` program, called `ufsboot`, must be

loaded on the boot disk by the bootblock program. This program must be placed in the boot area of the disk partition that will be booted. This bootblock program is different for each Sun platform. Copies of the program for the particular system can be found in `/usr/platform/<platform-name>/lib/fs/ufs`, where `<platform-name>` can be found using the `uname -i` command.

To install bootblock, perform the following commands:

```
# cd /usr/platform/`uname -i`/lib/fs/ufs
# /usr/sbin/installboot bootblk /
➔ dev/dsk/c0t3d0s0
```

Reboot the newly restored system


Once this last step is completed, the system can be halted and rebooted from the new

disk. Boot from the boot prompt and use boot's `-r` parameter to rebuild the device configurations:

```
OK boot -r
```

The system should boot up completely and be in the same state as when it was dumped.

Summary

Perhaps you've never had to restore your system, but sooner or later every serious system administrator faces this challenge. If you prepare for it and have the right information available, it will be an easy task. We hope this article gives you an idea of how to avoid pitfalls in the future. Use it as a reference when you're forced to restore your system. 



NET.UPDATE

Organizations that make the Internet work

by Edgar Danielyan

This is the first installment of net.update, a monthly column dedicated to Internet news and developments. In this issue, we'll discuss an overview of Internet governance and organizations that play vital roles in the Internet today and shaping of the Internet in the future. In future issues, we'll touch on both social and technical news, with brief commentaries from the author. We welcome your feedback and suggestions at edd@computer.org.

The Internet approach

The Internet has no central government, no parliament, and no police. Yet, it works, and it works much better than most governments. Let's take a look at the organizations that make the Internet work and at their internal structure.

Global organizations

Many internet organizations pull their members from all over the world. Here is an overview of these global organizations.

Internet Society

"To assure the beneficial, open evolution of the global Internet and its related internetworking technologies through leadership in standards, issues, and education."

Being a truly international and global membership organization, Internet Society (ISOC), found at www.isoc.org, has played a vital role in and for the Internet since 1992. It's a non-profit, non-governmental, open membership organization, with members from almost every country of the world. Its annual international conference, INET, draws hundreds of participants from all industries. It also organizes a number of network training workshops and tutorials, as well as publishes a monthly magazine, *On The Internet*.

Many chapters of the Internet Society are spread around the world and in many countries lead the Internet development and public awareness. Membership is available to

both individuals and organizations, and anyone is welcome to join the Internet Society.

Internet Assigned Numbers Authority

"Dedicated to preserving the central coordinating functions of the global Internet for the public good."

The Internet Assigned Numbers Authority (IANA), found at www.iana.org, assigns unique protocol, service, and port numbers, as well as coordinates the root name servers. It also publishes the *Internet Monthly Report* (IMR).

Internet Corporation for Assigned Numbers and Names

"The Internet Corporation for Assigned Names and Numbers (ICANN), found at www.icann.org, is the new non-profit corporation that was formed to take over responsibility for the IP address space allocation, protocol parameter assignment, domain name system management, and root name server system management functions now performed under U.S. Government contract by IANA and other entities. The Board of ICANN will be composed of nineteen Directors, nine At-Large Directors, nine to be nominated by Supporting Organizations, and the President/CEO (ex officio). The nine At-Large Directors of the Initial Board are serving one-year terms and will be succeeded by At-Large Directors elected by an at-large membership organization."

Internet Engineering Task Force

"The Internet Engineering Task Force (IETF) is a large open international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet. It is open to any interested individual."

The IETF, found at www.ietf.org, works on the technical aspects of the Internet. It consists of working groups (WGs) organized by topic into several areas, such as routing, security, applications, protocols, etc. The work is conducted through mailing lists; IETF also holds meetings three times a year. For more information on the internals of IETF, take a look at the Tao of the IETF at www.ietf.org/tao.html.

Internet Architecture Board

The Internet Architecture Board (IAB), found at www.iab.org, is the technical advisory body of the Internet Society. IAB appoints the

chair of the IETF, oversees the architecture of the protocols used in the Internet, as well as oversees the process of creation of Internet Standards. IAB also publishes *Requests for Comments* (RFCs); represents the Internet Society to other standards organizations; and advises the Internet Society on the technical, architectural, and procedural matters related to the Internet.

Internet Research Task Force

"To promote research of importance to the evolution of the future Internet by creating focused, long-term and small Research Groups working on topics related to Internet protocols, applications, architecture, and technology."

The Internet Research Task Force (IRTF), found at www.irtf.org, consists of a number of small, long-term research groups (RGs) which work on Internet protocols, applications, architecture, and technology. Participation in these research groups is by individuals, rather than organizations. The IRTF chair is appointed by the Internet Architecture Board (IAB). For more information about the Internet Research Task Force see RFC 2014.

Internet Engineering Steering Group

The Internet Engineering Steering Group (www.ietf.org/iesg.html) is composed of the area directors of IETF working groups. In addition, the General Area Director also serves as the chair of the IESG and of the IETF, and is an ex-officio member of the IAB.

Regional organizations

Many organizations have been formed that represent the interests of specific regions.

Reseaux IP Europeans (RIPE)

Seat: Amsterdam, The Netherlands

Areas served: Europe, part of Africa, part of Asia

The Reseaux IP Europeans, also known as European IP Networks or RIPE (www.ripe.net), Network Coordination Center (NCC) performs activities for the benefit of the Internet service providers (ISPs) in Europe and the surrounding areas. These are primarily activities that the ISPs need to organize as a group, although they may be competing with each other in other areas. The RIPE

NCC acts as the Regional Internet Registry (RIR) for Europe and surrounding areas.


Asia Pacific Network Information Center (APNIC)

Seat: Milton, Queensland, Australia
Areas served: Asia and Pacific

The Asia Pacific Network Information Center (APNIC) found at www.apnic.net, is a non-profit regional Internet registry organization for the Asia Pacific region.

American Registry for Internet Numbers (ARIN)

Seat: Chantilly, Virginia, USA
Areas served: Americas, Caribbean, and sub-Saharan Africa

The American Registry for Internet Numbers (ARIN) found at www.arin.net, is a non-profit organization established for administration and registration of IP network numbers and Autonomous System Numbers in the Americas, Caribbean, and sub-Saharan Africa. It replaced InterNIC IP/ASN registration services. 

Managing environment variables with Envy

by Rainer Dorsch

Recently, a user complained that our global path is so long that `tcsh` doesn't accept his personal extensions any more. Breaking down the big global configuration file into one file per application was an improvement, but not satisfying; users managed to load the environment for `jdk 1.1.1` and `jdk 1.2` concurrently. The result was that none of them worked as expected.

After these experiences, we were looking for a more robust solution for the management of our environments. We selected Envy, which satisfied our needs best. Envy consists of a collection of Perl and shell script developed by Joshua Nathaniel Pritikin, and distributed as free software under the terms of the Perl Artistic License. Envy makes managing different environments convenient and transparent.

In this article, we'll give an extended overview of Envy's features. We'll do so because one weakness of Envy is that the most useful part of the documentation isn't available before the program is completely installed.

Available environments

The `envy` command lists all available environments. **Listing A** contains an example output. A longer format with a description for all environments is also available.

A user can load an environment using the command `envy <environment>`. The command `envy jdk1-2` loads the environment for the Java development kit version 1.2. Installing an environment is completely reversible: `unvenv jdk1-2` uninstalls the environment for the Java development kit again. Although many people state initially that this feature is unnecessary, we found that it's quite useful.

For example, a user could have a tool called `prod` in his standard environment after login. The user works with `prod`, because it's very stable. But it has a bug, which forces him to sometimes use a newer beta version (`prod-beta`). In this case, he just types `envy prod-beta`. Envy unloads the environment for `prod` and loads the environment for `prod-beta` automatically.

Without Envy, the user has to remove `prod` by hand out of his environment. This is much more inconvenient and a common source of errors. Thus, Envy not only eases administration of complex environments, it also increases the productivity of users.

Each environment is defined in an ENV file. A detailed description is given in the section entitled "The ENV Configuration Files" later in this article.

Envy incorporates the concept of dimensions. Only one environment per dimension

can be loaded. This makes it possible to specify that it's impossible to load more than one version of a Java development kit at the same time.

The ENV files are written once and used by all shells through the same user interface. Many shells are supported, including sh, ksh, bash, tcsh, and csh. We use only bash and tcsh as login shells.

The installation

The only prerequisite for Envy is a working Perl 5.005 installation. A binary package of Perl in pkg format can be downloaded from www.sunfreeware.com.

Since Envy comes with installation instructions in the INSTALL file, we comment only on non-trivial parts (please do read the INSTALL file). We're running the bash shell when installing Envy. Set the installation directory of Envy with the code from [Listing B](#).

During the next login (best done via telnet), the links for the remaining shells are created. This pollutes your home directory with dot files; after that, you can change your login shell without rewriting configuration files for your environment. The command `envy` should now display all possible environments (only a test environment is installed by default).

Since the personal `.profile` and `.login` files are used for system purposes, new personal configuration files are necessary. These are located in the `.custom` directory: `profile` and `shrc` for bash and login and `shrc` for tcsh (all files without a leading dot).

At the end of `dot.profile`, `profile.part2` is sourced. This caused some difficulties at our site. `profile.part2` seems to be customized for Pritikin's site and probably needs some editing to satisfy your needs. The biggest problems were:

- **PS1 variable:** The prompt variable PS1 contains escape sequences in the default configuration, which sets the terminal title to the current host and path. This works fine with CDE, but crashes Openwindow's `cmdtool`. Setting PS1 to `PS1='[\u@\h] \w% '` is safe, but disables this feature.
- **PRINTER and LPDEST variables:** LPDEST and PRINTER are determined by


```
LPDEST=`ypmatch $HOSTNAME default_
printer`PRINTER=LPDEST
```

Listing A: Available environments

```
$ envy
  dailydb-dev          imag-dev-2509
x dev                 imag-dev-2513
  dev-area-setup      imag-dev-2519
  fame                imag-prod
  framemaker          jdk1-1-1
x fvwm95-2.0.43a     jdk1-1-6
  gcc                 x jdk1-2
  gems-1.9.1-dev      objstore
  gems-2.0-dev        openwin
  gems-2.0.alpha.14  printing
  gems-dev            prod
  gems-intl-rpts     x prod-new
  gems-prod           research
  imag-2519           reuters
  imag-dbo-testdb    solaris
```

Listing B: Code to set the installation directory of Envy

```
PERL5PREFIX=/usr/local/envy
export PERL5PREFIX
(for tcsh
setenv PERL5PREFIX=$HOME/test)
perl Makefile.PL
make test
make install

# a test
$PERL5PREFIX/bin/wrapper -s \
echo

mv $HOME/.profile $HOME/.profile.bak
ln -s \
$PERL5PREFIX/etc/dot.profile \
$HOME/.bash_profile
(for tcsh
ln -s \
$PERL5PREFIX/etc/dot.login \
$HOME/.login
```

This doesn't work for NIS+ environments. Since we didn't need this feature, we commented it out.

- **DISPLAY variable:** The DISPLAY guessing conflicts with ssh's handling of the DISPLAY variable. Commenting it out solved this problem.
- **Comments:** The comments printed are misleading after the above changes.

When the display manager *dtlogin* is in use at your site, you should customize the `$PERL5PREFIX/etc/login/.dtprofile` file. You should remove the comment (`#`) from `# DTSOURCEPROFILE=true` that all the applications you start from the desktop have the desired environment. Last but not least, you should change the welcome message in `$PERL5PREFIX/etc/envy/test.env` to something appropriate for your site. You can now access the online Help by entering *envy help* found in [Listing C](#).

Listing C: The *envy* online documentation.`$ envy help`

Envy 2.37 -- Environment Dimension Manager

Try:

```

envy help usage      for command line arguments
envy help custom    for a description of $HOME/.custom/ files
envy help author    for help writing .env files
envy help path      for an explanation of search paths
envy help env       for a list of envy specific environment
                   variables
envy help license   for licensing information

```

Send email to envy@listbox.com for support. Thanks!

Listing D: Documentation about the *.env* format`$ envy help author`

```

dimension java      # Declares dimension membership
desc Java 1.2 Test  # Description for 'envy list'
echo Java admin - call Joe x1212 # Outputs when loading
alpha              # Notify is alpha software
beta              # Notify is beta software
deprecated         # Notify is deprecated
error Java is no longer available. Sorry.
require Envy 2.16  # Uses required features from
                  2.16

require objstore   # Insures objstore is loaded
JAVA_HOME=/nw/prod/usr # Sets environment variable
JAVA_HOME:=$HOME/java # Overrides environment
                    variable
PATH+= $JAVA_HOME/bin # Prepend to colin separated
                    list
PATH+= $JAVA_HOME/bin # Append to colin separated
                    list
MYTOP=$ENVY_BASE   # Real path to .env file's tree
                    top
MYTOP=$ENVY_LINKBASE # Path to .env file's tree top

```

The ENV configuration files


A major task now is to convert environments for applications from a shell-specific format to Envy's format. Each environment is managed by an ENV file in `$PERL5PREFIX/etc/envy/` (global environments, available for all users) or `~xyz/.envy/` (personal environments, only available for the user *xyz*). The elements of an ENV file are described in the online documentation obtained by the command `envy help author` found in [Listing D](#).

The *dimension* key word ensures that conflicting environments aren't loaded at the same time. If there are several ENV files with the dimension *Java*, for example, `jdk1-1-1`, `jdk1-1-6`, and `jdk1-2`, `envy jdk1-2` ensures that `jdk1-1-1` and `jdk1-1-6` are unloaded (if necessary), before the environment for `jdk1-2` is loaded. *require objstore* ensures that `objstore` is loaded before the current environment. This is also reversible: unloading the current environment unloads `objstore` (except it's required by another loaded ENV file).

Final remarks

Although Envy makes the administration of the environment of many tools much easier, we found two shortcomings while converting scripts provided by the vendors of software tools:

- Scripts including aliases are difficult to convert, because Envy doesn't include alias management. At least reversibility in aliases is currently impossible with Envy.
- The documentation is very short. Furthermore, the most useful part is the online documentation, which isn't available before the installation is complete.

Another major concern with free software is always support. When evaluating the chances for support for free tools, we looked at the release history and the release interval of the tools. A long release history and frequent releases, indicate the author wants to give support as long as it helps him to improve the tool and he has the time to do it. The current version of Envy is 2.37, indicating a reasonable level of support. We submitted bug reports for Envy 2.30 and obtained fixes from the author within hours or, at most, two days. 

References

- Perl5.005 and many other freeware tools can be downloaded in the pkg format from <http://smc.vnet.net/>.
- The Perl Artistic License is available at www.perl.com/perl/misc/Artistic.html.
- The Envy download site is at <ftp://ftp.iguide.com/pub/mirrors/packages/perl/CPAN/authors/id/JPRIT/>.

Coming up...

- SNMP monitoring
- Web site configuration management
- Security issues

About our contributors

Edgar Danielyan spent some time studying US and UK law and is working as a network administrator and manager of a top-level domain of Armenia. He has worked for the United Nations, the ministry of defense, a national telco, a bank, and has been a partner in a law firm. He speaks four languages, likes good tea, and is a member of ACM, IEEE CS, SENIX, CIPS, ISOC, IPG, and many other much less known organizations. He can be reached at edd@computer.org.

Rainer Dorsch received his master in Physics in 1996 at the University of Ulm (Germany) and works at the University of Stuttgart (Germany) on his PhD in Computer Science. He started with UNIX in 1992, has worked with Linux since 1994, and has been heavily involved in Solaris administration since 1996. When he's not working, he likes hiking, cycling, or swimming in one of the beautiful lakes in the Alps or working in his orchard. He welcomes your comments. You may reach him via email at rainer.dorsch@informatik.uni-stuttgart.de.

Werner Klauser is an independent UNIX consultant working near Zurich, Switzerland. While not paragliding or roaring around on his Harley chopper, he can be reached by email at klauser@klauser.ch or on his Web page at www.klauser.ch.

H-W Schlote holds a diploma in physics. In 1996, he founded a software firm, SUFFIX Informatik GmbH, and is mainly working on projects in the UNIX environment. SUFFIX, for example, offers its own solution for load balancing and process distribution in heterogeneous UNIX clusters. He can be reached at H.Schlote@suffix.de.

Paul A. Watters is a research officer in the Department of Computing, Macquarie University, Australia. He can be reached at pwatters@mpce.mq.edu.au.

Inside Solaris (ISSN 1081-3314) is published monthly by
ZD Journals 500 Canal View Boulevard, Rochester, NY 14624.

Customer Relations

US toll free(800) 223-8720
Outside of the US(716) 240-7301
Customer Relations fax(716) 214-2386

For subscriptions, fulfillment questions, and requests for group subscriptions, address your letters to

ZD Journals Customer Relations
500 Canal View Boulevard
Rochester, NY 14623

Or contact Customer Relations via Internet email at zdjcr@zd.com.

Editorial

EditorGarrett Suhm

Assistant EditorJill Suhm
Managing EditorJoe Froehlich
Copy Editors.....Rachel Krayer

Christy Flanders
Taryn Chase
Contributing EditorsEdgar Danielyan
Rainer Dorsch
Werner Klauser
H-W Schlote
Paul A. Watters
Print DesignerMelissa Ribauda

You may address tips, special requests, and other correspondence to

The Editor, *Inside Solaris*
500 Canal View Boulevard
Rochester, NY 14623

Editorial Department fax(716) 214-2387

Or contact us via Internet email at sun@zsjournals.com.

Sorry, but due to the volume of mail we receive, we can't always promise a reply, although we do read every letter.

ZD Journals

General Manager..... Kelly Baptiste
Director of Customer Support Doug Neff
Director of Operations..... Kress Riley
Manager of Design and Production Charles V. Buechel
Manager of Product Marketing Mike Mayfield
Product Manager Jan Mater-Cavagnaro

Postmaster

Periodicals postage paid in Rochester, NY and additional mailing offices.

Postmaster: Send address changes to

Inside Solaris
P.O. Box 92880
Rochester, NY 14692

Copyright

Copyright © 1999, ZD Inc. ZD Journals and the ZD Journals logo are trademarks of ZD Inc. *Inside Solaris* is an independently produced publication of ZD Journals. All rights reserved. Reproduction in whole or in part in any form or medium without express written permission of ZD Inc. is prohibited. ZD Journals reserves the right, with respect to submissions, to revise, republish, and authorize its readers to use the tips submitted for personal and commercial use. For reprint information, please contact Copyright Clearing Center, (978) 750-8400.

Inside Solaris is a trademark of ZD Inc. Sun, Sun Microsystems, the Sun logo, SunSoft, the SunSoft logo, Solaris, SunOS, SunInstall, OpenBoot, OpenWindows, DeskSet, ONC, and NFS are trademarks or registered trademarks of Sun Microsystems, Inc. Other brand and product names are trademarks or registered trademarks of their respective companies.

Printed in the USA.

Price

Domestic\$99/yr (\$9.00 each)
Outside US\$119/yr (\$11.00 each)

Our Canadian GST# is: R140496720. CPM# is: 1446703.

Back Issues

To order a back issue from the last six months, call Customer Relations at (800) 223-8720. Back issues cost \$9.00 each, \$11.00 outside the US. You can pay with MasterCard, VISA, Discover, or American Express. Collections of back issues are available on CD-ROM as well. Please call for more information.

Are you moving?

If you've moved recently or you're planning to move, you can guarantee uninterrupted service on your subscription by calling us at (800) 223-8720 and giving us your new address. Or you can fax us your label with the appropriate changes at (716) 214-2386. Our Customer Relations department is also available via email at zdjcr@zd.com.

PERIODICALS MAIL

Sun Technical Support
(800) 786-7638



*****3-DIGIT 480


C: 7661905 00002096 04/00

23

57

CLINTON TOWNSHIP MI 48035-4218

Please include account number from label with any correspondence.



Fixing network settings

by Paul A. Watters

I came into work this morning to discover that one of my colleagues had been fiddling with the network settings on my Sparc. Now I can't see any machines on my local subnet, or connect to remote hosts. What are the first things to check?

The first thing to check is that your computer is physically connected to the nearest network hub, and that your network card status LEDs indicate activity. If this is the case, then check that the loopback connection is working, and your network identification information is correct with the following:

```
ping 127.0.0.1
```

If you get a message like

```
127.0.0.1 is alive
```

then you're on the road to recovery. If you

can ping other machines on the local subnet and remote hosts by using IP addresses, but not DNS names (assuming that NIS or NIS+ isn't being used), then checking the DNS configuration is the next obvious step. First, verify that `/etc/resolv.conf` contains the line

```
hosts: files dns
```

which determines the order of resolution as local host files (typically `/etc/hosts`), which should contain definitions for localhost identity and the actual DNS name and IP address:

```
127.0.0.1    localhost
10.20.30.40  phoenix loghost
```

If this all seems in order, check `/etc/resolv.conf` to verify that a nameserver and domain are specified, that is:

```
domain arizona.com
nameserver 10.20.30.128
```

Resizing partitions

by Paul A. Watters

One of my partitions is getting very full. Is there any way to resize the partitions without risk, or rationalize the space somehow?

Although there are products in the marketplace which claim to be able to resize partitions online, it's always a risky business, and there's a chance that once the partitions have been resized, that data on partitions whose size is reduced will be unavailable. Of course, with a sensible backup policy in place, this isn't really an issue.

One possibility is to take a full dump of each partition's data on tape (using `ufsdump`),

resize the partitions using `format`, and then restore the data to the resized partitions. Clearly, this strategy only works if the resized partitions have sufficient disk capacity to cope with the data being reloaded. Another possibility for the less adventurous is to tar up the contents of one filesystem, and then untar it directly onto a larger slice (for example, when a new disk has been acquired). An example command if we wanted to copy the entire contents of `/u01` to `/u02` would be:

```
$ cd /u01; tar cf - . | (cd /u02;
➤ tar xfbp -)
```