

TOPICS

An OS/390 Newsletter for Release 9

February 2000

Our second issue ...

Welcome to the second edition of our *Hot Topics Newsletter!* We enjoyed putting the first one together, and are gratified to learn that many of you have found it useful. That's the fundamental objective. Don't stop letting us know what you think. We check our e-mail at least once a day (on weekdays), and will get back to you as soon as we can. We welcome comments both about the Newsletter itself and about any aspect of OS/390®. Any comment or suggestion about the style or the contents of the Newsletter, we will personally address. Any comment about the system, we will forward to the appropriate expert.

Some of your mail asked how you could get on our mailing list. The short answer (now that the System Library Subscription Service has gone away) is that there is no way. But if you send us an e-mail, we'll send you a note twice a year, reminding you about all the places where issues are available (see page 31 of this issue).

Finally, because we haven't had new office furniture in years, our desk chairs made the usual noises one afternoon as we leaned back to discuss various articles. Hence the name of a little character we introduce with this issue.

The Editors,
newsletr@us.ibm.com

UNIX on OS/390: A powerful blend

This issue also includes a supplement on various aspects of UNIX® on OS/390. It attempts to make UNIX and MVS™ mutually intelligible to those whose primary operating system background is one or the other. An excellent source of help for OS/390 UNIX is the MVS-OE mailing list. See the supplement for information about how to subscribe.

This 16-page supplement follows page 32 of this newsletter.

Demystifying TCP/IP

Again and again, we've heard that the Internet protocol, TCP/IP, has been a source of confusion to you. To clear the air, this issue features several articles that address TCP/IP topics that you've



Paul Rowntree

identified as especially mystifying. Here's hoping they help you.

The rebirth of TCP/IP on OS/390 • pg 6
Useful IP sites on the Web • pg 7
Clearing up the IP migration mist • pg 8
Concerned about migrating to TCP/IP in OS/390 V2R5 or later? • pg 9-11
Choosing between SNA and IP integration options • pg 12-13
Dynamic VIPA? No problem – with a little planning • pg 14-15
TCP/IP application workload balancing in the S/390 Parallel Sysplex® • pg 16-17
Improved TCP/IP performance metrics • pg 18-19
A roundup of IP security • pg 20-23
Is OS/390 TCP/IP ready for DB2®? • pg 24-27

Find out What's Hot!

OS/390 Release 9 documentation — streamlining book shipments • pg 2-3

OS/390 — What's new in R9 • pg 3

Take the R out of SURF and what have you got? S/390 service on the Web! • pg 4

How to print PostScript, PDF, and SAP data on AFP printers with Infoprint Server • pg 5

Introducing LOOKAT, a prototype message help facility for OS/390 • pg 28

The growing wizardry of OS/390 • pg 29

Easing your way through those large OS/390 collection kits • pg 30-31

Here's where you can find Hot Topics • pg 31

The URLy bird gets the info • pg 32

OS/390 Release 9 documentation — streamlining book shipments

by Paul Gorgen

Over the last several years, we've made a number of viewing and printing enhancements to our softcopy information. Also, better tools to help you manage your softcopy repositories have made you rely less on IBM®-printed hardcopy books. If you are now using softcopy primarily and encouraging all users you support to do likewise, the thought: "I wish I could trade these hardcopy books that come with my order for something else" has certainly crossed your mind.

With OS/390 Release 9, we've fine-tuned the list of books that you automatically receive. You will automatically receive the *OS/390 Installation Planning Kit* just as you always have. The *OS/390 Installation Planning Kit* is a white and blue box containing:

- *OS/390 Introduction and Release Guide*
- *OS/390 Planning for Installation*
- *OS/390 Information Roadmap*
- *OS/390 Hot Topics Newsletter* (what you are reading now)

You'll also automatically receive a key set of 36 additional hardcopy books. This set includes a number of installation manuals, as well as some of the most-used books such as *MVS System Messages*.

To see a list of these books, go to the IBM Softcopy web page at <http://www.ibm.com/s390/softcopy>.

If you want to purchase additional hardcopy books with your order, an optional feature code (8007) allows you to order about 50 more. This set of books includes:

- Any OS/390 books with messages and codes that are not already included in the set of 39 books that you automatically receive with your order.
- The rest of the books that used to ship automatically with OS/390 Releases 6 through 8, and that are now not part of the basic set. To see a list of these books, you can go to the IBM Softcopy web page at <http://www.ibm.com/s390/softcopy>.

There also continues to be the option to purchase subscriptions to the *Redbook CD-ROM* collection and the *Security Server CD-ROM* collection with your OS/390 product order.

So, Release 9 has a streamlined hardcopy distribution plan. For those of you who want all the OS/390 message and codes books in hardcopy, it is now much simpler. Instead of having them scattered in the previous huge feature code for the base element books (which might have required you to buy more hardcopy books than you really wanted) and throughout the small feature codes we used to have for each optional feature, they are now all in two places (in the hardcopy you automatically receive and in the 8007 feature code).

And, as always, don't forget to continue to check out the OS/390 Web site (<http://www.ibm.com/s390/os390>), and click on the Library, for the latest online books, PDF files, wizards, and other information.

OS/390 — What's new in R9

by Dick Wagenaar

The marketplace transformation from traditional on-line transaction processing to e-business transaction processing is continuing to challenge business everywhere.

But with the right solutions, the opportunities are tremendous. Release 9 of OS/390 continues the tradition of providing you with protection of your past investments, while offering a platform uniquely suited to help you take full advantage of e-business.

We offer new security functions, for example, to help you go after more e-business. The OS/390 Open Cryptographic Services Facility has functional enhancements in the handling of OCSF Cryptographic Contexts, and the base now includes some functions that were in former features. The System Secure Sockets Layer includes two new application interfaces.

Language Environment® has some exciting new enhancements in reliability, availability and serviceability.

OS/390 C/C++ has introduced many new compiler options and sub-options and #pragma directives.

OS/390 UNIX System Services include enhanced systems management capabilities, improved debugging tools, and additional common UNIX functions that ease the porting of applications from other platforms.

Server Consolidation is helped by Distributed File Service, HCD, HCM, and DCE enhancements. HCD, for example, now has support for LDAP exploitation, and HCM now gives you the facility to use the GUI of the ESCON® Manager or I/O Operations of System Automation for OS/390, and there are physical connectivity enhancements. DCE has an enhancement that eliminates a potential bottleneck by allowing the credentials cache to be shared within the same sysplex.

WLM introduces the concept of multisystem enclaves, which now allows you to group work and have it be managed as a unit across all systems in a sysplex. And JES3 has some improvements in the specifying of NJE node names and DEVICE statements. Finally the BCP now monitors structures and tells you when they're getting full, and includes trace enhancements for transactions.

Need education?

IBM Learning Services offers IBM customers dozens of OS/390 courses ... both instructor-led and distance learning.

Take a look at our Web site at <http://www.ibm.com/services/learning/spotlight/s390.html> for more information. We also offer the OS/390 Update Information Service, which is a subscription service that gives you twice-yearly technical updates on each new OS/390 release to help you make an informed migration decision. Subscribers receive this training via CD-ROM or over the Internet. OS/390 experts are available via e-mail to answer specific training questions, and frequently asked questions are summarized in a FAX data base. Here's the Web page where you can sample this new way to stay current on OS/390:

<http://www.ibm.com/services/learning/demo/s390updt.html>



Stacy Newman

Take the R out of SURF and what have you got? S/390 service on the Web!

by Ted Posner

All you S/390 system programmers out there surfin' the Web! There's a **big wave** headed your way. Ride it. It's called the S/390 Service Update Facility (SUF). What is it? Customized OS/390 service deliverables! Which means that now you can order, take delivery, and even automate initial installation steps for OS/390 maintenance right over the Web.

We've got three levels of service ... all available over the Web.

You can **hang ten** for "preventive service." With this level of service, you order, receive and install OS/390 Recommended Service Upgrades (RSUs). With these OS/390 service deliverables, you update one or more of your Target Zones to a new service level. Using SUF, you get just the maintenance you require, no more, no less. When you add the automation of SMP/E RECEIVE to this, SUF is the fastest, easiest way to get an RSU.

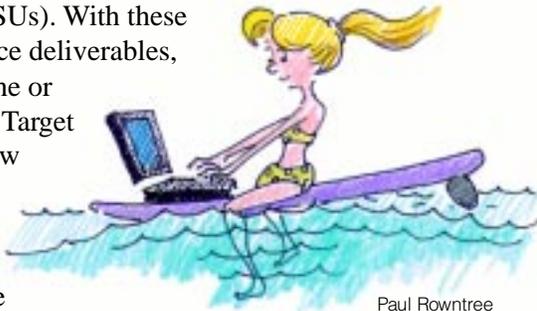
You can **hang five** for "critical service." With this level, you order, receive and install Hiper, PE, and Y2K fixes. You use SUF to order all the Hipers, PEs, and Y2K fixes that you are missing for one or more of your target zones. Again, SUF automates the RECEIVE.

You can **hang one or two** for "corrective service." With this level, you order, receive and install one or more PTFs (ordered by PTF or APAR number). Again, SUF delivers just what you order, no more, no less. SUF ensures that any missing pre- or co-requisite PTFs are also delivered. This cuts down your research time, and the time you spend waiting for additional requisites.

In addition to delivering OS/390 service, you can use SUF to order and receive VM and VSE service.

For more information about SUF, put the R back in and surf on over to our homepage at:

<http://www.ibm.com/s390/suf/>



Paul Rowntree

How to print PostScript, PDF, and SAP data on AFP printers with Infoprint Server

by Allida Shoning

Do you want to print documents generated by PC-based or OS/390 UNIX applications on IBM high-speed Advanced Function Presentation (AFP) printers such as the IBM Infoprint 60? Now, with Infoprint® Server and Infoprint Server Transforms, you can automatically transform PostScript, PDF, SAP R/3, and PCL documents into AFP format on your OS/390 system before printing. These are the two products you need:

Infoprint Server for OS/390, a feature of OS/390 (V2R8, or higher), lets you print either from the local OS/390 system or from remote platforms such as Windows™, UNIX, AIX®, and OS/2®. Infoprint Server automatically invokes transforms on the OS/390 system to convert data from one printer format to another.

Infoprint Server Transforms for OS/390, a separate product that is currently offered at no additional charge to Infoprint Server customers, provides these OS/390 transforms:

PCL to AFP
PDF to AFP
PostScript to AFP
SAP R/3 ABAP to AFP
SAP R/3 OTF to AFP

Later releases of Infoprint Server Transforms will provide additional transforms. And customers can provide their own datastream transforms by using the transform API.

Here's a typical printing scenario:

An installation wants to print three copies of a 400 page OS/390 publication on a high-speed Infoprint 60 printer. These publications are in PDF format in an hierarchical file system (HFS) on the OS/390 system.

To automate printing, the OS/390 administrator uses ISPF panels to define the Infoprint 60 in the Infoprint Server Printer Inventory. The administrator assigns a name to the printer, for example PRT60, and requests that the PDF to AFP transform (provided with Infoprint Server Transforms) be used to convert PDF data streams to AFP format. The administrator might also specify default printing options, such as printing on both sides of the paper.

To print the publication:

1. The job submitter logs onto OS/390 UNIX System Services by using the rlogin function on a UNIX system.
2. Using the OS/390 UNIX System Services printing command (lp) provided with Infoprint Server, the job submitter requests that three copies of the file be printed on printer PRT60.
3. Infoprint Server automatically calls Infoprint Server Transforms to convert the PDF file to AFP format.
4. Infoprint Server creates an output data set on the JES spool, with output parameters that specify three copies and duplex printing.
5. PSF for OS/390, which is a separate product, selects the output data set from the JES spool and prints three copies, on both sides of the paper, on the Infoprint 60 printer.

The rebirth of TCP/IP on OS/390

by Alfred Christensen

Alfred Christensen is currently a senior systems developer in the CS OS/390 strategy and design group. He has 25 years of experience with IBM mainframe software and is an IBM-certified consulting IT architect in network computing solutions.

OS/390 V2R5 provided a new TCP/IP component in the Communications Server for OS/390 to completely replace the old TCP/IP for MVS V3 product. Initially ported from VM, TCP/IP for MVS focused on IP network access and base server and client functions, including the ability to write your own MVS socket programs.

As the commercial importance of TCP/IP on MVS and OS/390 grew, so did customer demand for more function, higher availability, improved reliability, and better performance. A key requirement was the ability to scale socket workload across multiple processors and even OS/390 images in a Parallel Sysplex. The original code was based on a single task in the TCP/IP address space. That is, work performed in the TCP/IP protocol stack took place under a single MVS task control block (TCB) that could be dispatched only on one S/390 processor at a time. Within an n-way S/390 server, the TCP/IP workload was limited to the capacity of a single processor! A new product was needed.

The new TCP/IP protocol stack allows multiple application programs to read or write data from (or to) the TCP/IP networks simultaneously, using all available processors in an n-way S/390 server.

We attempted to keep all existing external functions and APIs intact. One external API was removed — the low-level assembler API directly on top of the IUCV/VMCF subsystem. Software vendors, as well as MVS TCP/IP users, were notified 3 years in advance that the API would disappear and most ISVs converted their applications to one of the other socket APIs. Existing applications such as traditional TCP/IP V3R2 C-socket programs,



Paul Rowntree

sockets extended programs (callable or assembler macro), ONC/RPC, XTI, X-Windows, interpreted REXX socket programs, or Pascal programs run unchanged on OS/390 V2R5 (or higher). Application developers do not need to re-code them, change any external configuration files they have created (such as TCPIP.DATA), or learn about OS/390 UNIX. Depending on your current level of TCP/IP for MVS, some programs can be re-linked to achieve optimal performance, but that's the extent of the work that needs to be done.

Performance and scalability were top TCP/IP for OS/390 priorities. OS/390's future involves enabling popular UNIX applications — such as Web services, DCE-based

services, and directory services — so we decided to take advantage of “UNIX Sockets” for these applications. To meet the performance objectives of such applications, the code path between these applications and the TCP/IP stack had to be optimized, and therefore the TCP/IP stack and UNIX System Services became much more integrated.

You do not need to be a UNIX expert to run the new TCP/IP stack! After the stack is initialized, you can move your existing MVS-based socket workload to the new TCP/IP without any major concern for UNIX. Should users require new OS/390 server functions, you may have to learn more about the OS/390 UNIX environment. OS/390 UNIX allows programs written to the open standards as defined by POSIX and XPG4 to run on OS/390 along with traditional MVS programs. Instead of developing a new DNS server for OS/390 from scratch, we used standard DNS/Bind source code to run as a standard DNS server on OS/390. Other standard components such as Sendmail, the TFTP (Trivial File Transfer Protocol) server, and the DHCP (Dynamic Host Configuration Protocol) server were also ported from UNIX and they behave somewhat as “normal” UNIX programs do. For a migration from a standard TCP/IP for MVS V3R1 or V3R2 level, you wouldn't initially need any of these new servers. In some situations, new TCP/IP functions will require that you familiarize yourself with OS/390 UNIX. But the benefit is that you will be able to offer users a completely new set of OS/390-based services, often at a price significantly lower than what they currently pay for running the same types of services on various UNIX boxes.

Useful IP sites on the Web

by Michele Carlo

Michele Carlo has been an IBM Information Developer for 18 years and is currently engaged in revising IP information within the SecureWay Communication Server library to better meet customer needs.

Here are some Web sites where there's more SecureWay Communications Server information:

All about IP and SNA

Find out about IBM's SecureWay® Communications Server product at: <http://www.software.ibm.com/network/commserver/>

Learn about tasks, services, and information for Communications Server for OS/390 by accessing: <http://www.s390.ibm.com/nc/techinfo/comserve.html>

Receive guidance on the Communications Server product set and help migrate to new Communications Server products and architecture by viewing the IBM Communications Server Software Roadmap at: <http://www-4.ibm.com/software/network/technology/roadmap/>

Interested in IP Performance?

See: http://www-4.ibm.com/software/network/commserver/library/whitepapers/csos390_v2r6perf.html

Need general information on OS/390 tuning/performance? It's available at: <http://www.s390.ibm.com/products/oe/bpxa1tun.html>

Have an IP problem?

Try IBM SecureWay Communications Server support at: <http://www.software.ibm.com/network/commserver/support/>

Want some service tips?

These tips are generated from customer problems to help you avoid the same ones. See SecureWay Communication Server Service Tips at: <http://www2.software.ibm.com/eNetwork/Tips.nsf/Tips>



Paul Rowntree

Want to read IP redbooks?

See IBM Systems Center publications at: <http://www.ibm.com/redbooks>. Review the ITSO redbooks for CS V2R5 and V2R6 installation and usage (specifically SG24-5227, SG24-5228, and SG24-5229.) Enter the manual numbers in the search window. To review the ITSO redbook on using TCP/IP in a Sysplex, see SG24-5235.

How about some hot flashes?

Think IBM Systems Center flashes at: <http://www.ibm.com/support/techdocs> under FLASHES.

To get a specific flash, select View by number, then search for: **W98019**, Communications Server for OS/390 V2R5 TCP/IP Migration Tips, or **W98042**, Communications Server for OS/390 V2R6 TCP/IP Migration Tips, or **W3192**, Communications Server for OS/390 V2R7 TCP/IP Migration Tips, or **W3196**, Communications Server for OS/390 V2R8 TCP/IP Migration Tips, or **N3190**, a Washington System Center Flash on Multicast and VIPA Considerations When Configuring OMPROUTE on OS/390 TCP/IP, or search for **N3191**, a Washington System Center Flash on Multicast and (O)Routed on MVS and OS/390 TCP/IP

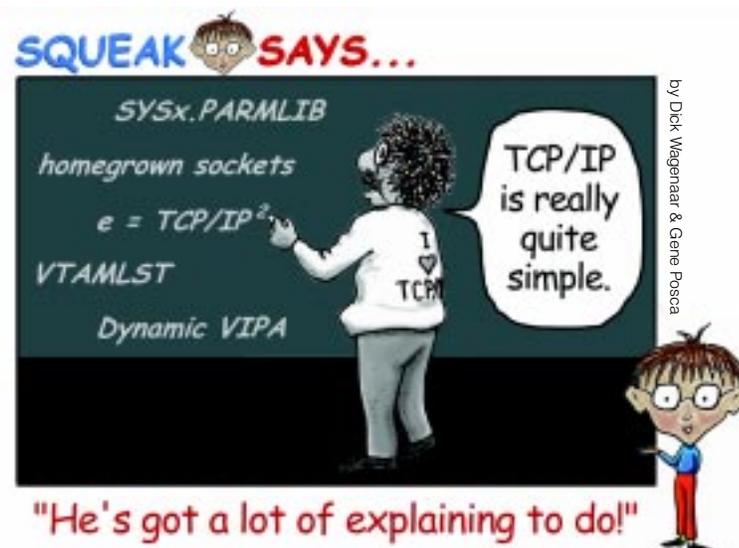
Need historical perspective?

VTAM and TCP/IP are explained at: <http://www.software.ibm.com/network/commserver/about/csos390.html>

Want to read about TCP/IP for MVS?

See: http://www-4.ibm.com/software/network/commserver/about/find_tcpipmvs.html

Find background information on the redesigned Communications Server IP stack at: <http://www.s390.ibm.com/oe/bpxqt10.html>



"He's got a lot of explaining to do!"

Clearing up the IP migration mist

by Michele Carlo

If you've been frustrated making required changes when going from one release of TCP/IP to the next, you'll be glad to know we've improved the IP migration documentation. We listened to you and used much better formats.

Take a look at the latest version of *SecureWay Communications Server: IP Migration* (SC31-8512-04), and see what we mean.

A note on migration:

All OS/390 migration information is being structured like this. Now you can migrate and worry less about important information you might overlook.

Before V2R8, the *IP Planning and Migration* guide was pretty unstructured. Lack of clarity contributed to less-than-ideal migration experiences. Our new book, *IP Migration*, has a very structured format.

A new roadmap links major functional enhancements to releases; a special section shows how interfaces have changed across releases; and most important of all, complete information about migration tasks for every new function appears in simple tables, including lists of required and optional tasks.

Migration information about major applications, such as FTP, Telnet, SNMP and DNS remains in separate chapters, because that is what you prefer. But you'll see more order and the easier-to-understand format in those chapters, too.



Paul Rowntree

Concerned about migrating to TCP/IP in OS/390 V2R5 or later?

by Alfred Christensen

Uncharted territory is always scary. Moving to TCP/IP in OS/390 V2R5 (or higher) from TCP/IP for MVS V3R1 or V3R2 is like entering the unknown. But the rewards – performance, reliability and many new functions – are worth the effort. V2R5 TCP/IP enhancements make use of many new OS/390 technologies, especially those offered by OS/390 UNIX System Services. You've never used UNIX System Services? Be sure to familiarize yourself with the basics. You might start with the article on the re-birth of TCP/IP in this issue. Then take up the *UNIX System Services Planning Guide*. It provides the information you need to initialize the UNIX environment and to migrate to the new TCP/IP.

TCP/IP is part of Communications Server for OS/390 and uses functions that are common to VTAM® and TCP/IP. A new Communications Storage Manager (CSM) and a common Data Link Control (DLC) allows sharing of certain network interfaces between TCP/IP and VTAM. If your background has been primarily TCP/IP, you need to either refresh your VTAM administration and configuration skills, or get help completing the proper TCP/IP network interface VTAMLST definitions.

Before you start migration activities, please read the *V2R8 IP Planning and Migration Guide*. Your time will be well spent.

Things to know up front

You need to know a few things before you start your migration.

- The new TCP/IP doesn't support assembler IUCV API sockets, supported until TCP/IP V3R2. So you need to identify the socket applications that use this API. Since this API was entirely assembler-based, it is unlikely that your own socket applications were developed using it, but some third-party products did. Most have provided newer versions that use socket APIs supported by the new TCP/IP. So if you have such a product, plan for an upgrade before you migrate to the new TCP/IP. If you have homegrown socket applications that use the IUCV API, convert them to use either the sockets extended assembler macro API or the UNIX System Services callable BPX sockets API. Plan this conversion well in advance of migrating to the new TCP/IP.

- The new TCP/IP doesn't support offload devices. In the past, offload of certain TCP/IP functions to an outboard channel-attached offload box reduced S/390 CPU usage. The new TCP/IP is so much more efficient than the old that S/390 CPU offload is not necessary. Offload devices, such as an IBM 3172, can be reconfigured to act either as LCS devices or as channel-attached router devices. For more planning information, please refer to the IP URLs article elsewhere in this issue.

Since OS/390 V2R8, the ServerPac installation and customization procedure helps define and start an IVP (Installation Verification Procedure) for the TCP/IP environment. You may be able to use some of the IVP sample definitions or sample jobs as a reference.

Set up UNIX System Services

Before installing TCP/IP, you need to establish a full-function mode UNIX System Services environment, including a root Hierarchical File System (HFS). The root HFS is fully pre-built as a large MVS data set during ServerPac install and ready to use. UNIX System Services are initialized automatically during IPL of OS/390. Currently, HFS data sets have to be allocated on SMS-managed DASD volumes. When UNIX System Services initializes, the UNIX configuration options are read from the BPXPRMxx member in your SYSx.PARMLIB data set. Before starting TCP/IP, you and your MVS systems programmer need to review some of those definitions to add specifications that define how UNIX System Services communicate with your TCP/IP system address space. Include these definitions in FILESYSTYPE and NETWORK statements. Use either the INET approach or the common INET (CINET) approach. INET allows one TCP/IP system address space to be used with UNIX System Services. CINET allows multiple TCP/IP system address spaces, including AnyNet, to be used with UNIX System Services. We strongly recommend you use INET. You generally should not start more than one TCP/IP system address space on an OS/390 system.

Set up security

Beginning with OS/390 V2R5, many standard TCP/IP functions — including all S/390 sockets programs — use UNIX functions. To use these functions, user and group IDs must have UNIX System Services definitions in a security database segment called an OMVS segment. RACF® supplies a utility program that assigns OMVS segments to all your user and group profiles. You might also use what is known as a default OMVS segment. If a user with no OMVS segment defined tries to use UNIX functions, the system uses the default OMVS definitions you defined and enabled. However, before you use the default OMVS segment support, consider implications to your installation's security policy. The user ID associated with the TCP/IP system address space must be defined with an OMVS segment that has a UID (UNIX ID) of zero (a UNIX super-user). User IDs with zero specified must be associated with servers such as the FTP server, the OMPROUTE server, the DNS server, and the SNMPD agent.

TCP/IP V2R5 uses standard Security Access Facility (SAF) security authorization for many functions, such as operator command support. You may want to add OPERCMDS resources to protect the various operator commands that TCP/IP allows. UNIX System Services allows detailed control over which programs are allowed to request an ID and password from an end-user and then execute under that user ID. Servers like the FTP server need this ability. To enable this level of control, define a BPX.DAEMON facility class resource in your security product. Give the FTP server started task user ID READ access to the facility class resource and define a few selected load libraries as program controlled. Although this may sound a bit complicated, detailed directions in the *IP Planning and Migration Guide* are straightforward and quite easy for a security administrator to implement.

Set up SyslogD

Many operating systems use a special programming interface to write trace information or error messages to a common location. This common location on UNIX systems is the SyslogD server. If you do not start the SyslogD server on OS/390, many of these messages will be directed to your MVS log. We suggest you configure SyslogD to store the messages it receives into files in the HFS. You can create the SyslogD configuration and start the SyslogD server through a standard MVS JCL procedure. Start SyslogD before you start the TCP/IP system address space.

Prepare to start TCP/IP

Replace your old TCP/IP started task JCL procedure with a new sample (TCPIPROC in tcpip.SEZAINST). You can use your old TCP/IP PROFILE without modifications. Because some of the definitions are no longer valid, you will receive warning messages when you start the new TCP/IP. During a period where you may need to use the same TCP/IP PROFILE with both an old TCP/IP and the new TCP/IP, you can just ignore those messages. Later, clean up your PROFILE and start exploring some of the new configuration statements, such as IPCONFIG, TCPCONFIG, and UDPCONFIG. Because the new TCP/IP shares Data Link Control (DLC) functions with VTAM, you need to have VTAM running for TCP/IP to be able to communicate through its network interfaces. Start VTAM before TCP/IP.



Paul Rowntree

Prepare to start the servers

Here are some ideas on what to do for some servers and what to look out for. The new TCP/IP TN3270 server supports the TN3270e protocols including 3270-printer support. The protocol specifies that if either the client or server is TN3270E capable, it should first try TN3270e and then back off to the older TN3270 protocol. Some clients don't behave according to protocol specifications, fail, and break the connection as soon as the OS/390 TN3270 server tries to negotiate TN3270e. To support these clients, the OS/390 TN3270 server has a configuration option that says NOTN3270E to prevent it from trying to negotiate TN3270e. With OS/390 (V2R6, or higher) you can configure multiple TN3270 server ports individually.

You could configure an alternate port as NOTN3270E and instruct users with old client software to connect there instead of the default TN3270e port number. Conduct tests with all your TN3270 client software and the new OS/390 TN3270 server. The new FTP server supports access to both MVS data sets and HFS files. The tasking model is quite different from the old TCP/IP V3R2 FTP server that processed all client connections as subtasks within a single address space. The new FTP server uses a separate address space for every client that connects to the server,

virtually removing any limitations on the number of concurrent FTP sessions — the limit is the number of concurrent address spaces your configuration supports. Since OS/390 V2R6, TCP/IP supports RIP V1, RIP V2, and OSPF with a new routing server called OMPROUTE, the old OROUTED server is still supported, but only for RIP V1 and RIP V2.

This article can't provide you with every migration detail, but I hope it makes you feel more confident about TCP/IP migration. There are many new things to learn, but the value to your OS/390 systems users is great. The good news is that when you've jumped over this TCP/IP migration hurdle, your next OS/390 upgrades will be so much easier.

Shhhhh! This is a library.

But here's something we want to shout from the rooftops!

Get your own personal librarian

Manage your library
of BookManager
books with the
SoftCopy Librarian.

No more uploading books one at a time.
No need to do it by hand.
No more lists of what's in your library.
No more searching through the CD catalog to
identify books that are new or have changed.

It's the future!

Read about the SoftCopy Librarian in any CD collection, and to download the Softcopy Librarian package visit our Web site at:

<ftp://ftp.software.ibm.com/ps/products/ibmreader/tools>

Choosing between SNA and IP integration options

by Richard Tobacco

Richard Tobacco is the brand manager for Communications Server for OS/390. Richard has supported networking products for nearly 15 years.

Most businesses rely on legacy applications residing on S/390. Systems Network Architecture (SNA) has traditionally provided network access to these applications and databases. But the application development environment has drastically changed. It is more and more aimed at e-business over TCP/IP networks. Initially, businesses created and supported these SNA and TCP/IP networks separately. Today's cost-conscious network managers seek ways to consolidate their SNA traffic onto the IP network. Since rewriting an estimated three billion-dollar investment in legacy SNA-based applications to IP is clearly not justified, transport solutions are required. Let's talk about a few.

The IP client solution - TN3270 emulation

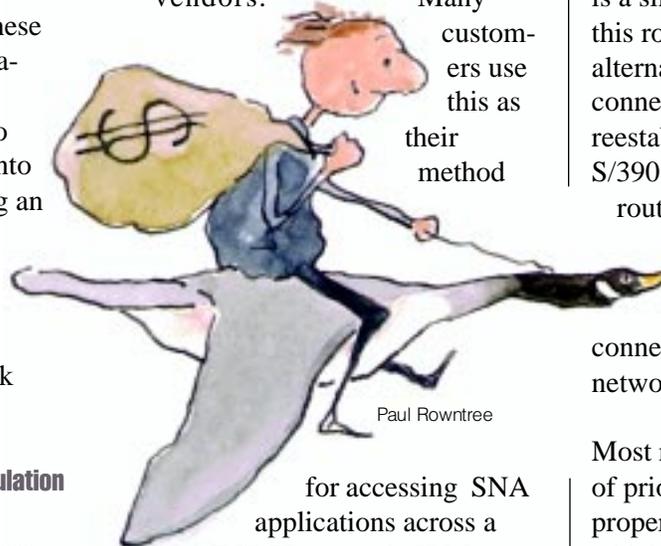
Today, many businesses use TN3270 as their method of accessing SNA applications by converting end-user terminals to be TCP/IP-based. Such terminal emulation allows users to access Web-based applications, new applications written for TCP/IP, and existing SNA applications. The end-user emulators are TN3270 clients that talk to TN3270 servers across the IP network. Servers can be outboard (a separate device) or contained within S/390. S/390 can support more than 64,000 TN3270e sessions.

Solutions for the SNA client

It is often not possible or too costly to convert all end devices to TCP/IP. These devices could be banking terminals requiring a SNA communications protocol (for example, SDLC) or they could be actual 3270-like devices coax-attached to a 3274-like terminal controller. In these situations, businesses can still merge onto an IP network for transport of SNA traffic, if the SNA traffic is first encapsulated into IP.

In 1992, IBM introduced Data Link Switching (DLSw) as a means for transporting SNA data across a TCP/IP network. DLSw gained widespread acceptance by routing vendors.

Many customers use this as their method



for accessing SNA applications across a TCP/IP network. IBM

created Enterprise Extender (EE) as an open alternative way to integrate SNA applications onto an IP network. Cisco Systems, Inc. routers also provide the same function, called SNA Switching Services (SNASw).

TCP encapsulation - DLSw

DLSw is the encapsulation of SNA protocols within a TCP/IP 'package'. The package is a useable TCP/IP address known to DLSw-capable routers. At both the sending and receiving locations, DLSw-capable routers terminate the SNA connection and locally acknowledge transmissions. Connection setup and maintenance is process-intensive involving link-level acknowledgement, TCP retransmit, congestion control, protocol translation and data store-and-forward. This is a significant router burden even for powerful routers.

With DLSw, the data-center router is a single point of failure. Should this router fail, even though an alternate path may exist, all SNA connections would have to be reestablished. Traffic between S/390 server and the data-center router, where the DLSw session is terminated, is SNA. Therefore, the S/390 server supports both IP and SNA connections, even though the network is only IP.

Most routers provide some form of prioritized queuing; however, properly identifying the priority at which a SNA packet should be sent has traditionally been difficult. With DLSw, traffic prioritization is handled on a link basis and multiple links must be defined to the same SNA device.

UDP encapsulation - Enterprise Extender

Enterprise Extender (EE) is an extension to High Performance Routing, a lower overhead, less complex SNA. It provides efficient encapsulation of SNA application traffic within UDP frames by EE-capable devices at the edges of an IP network. To the IP network, the SNA traffic is UDP data-grams that get routed without hardware or software changes to the IP backbone. To the user, the session is 'normal SNA' with predictable performance and high availability. There is no protocol transformation. Encapsulation is performed at the routing layers enabling efficient use of the IP infrastructure for accessing SNA-based data.

Advantages of Enterprise Extender

EE is available on Communications Server for OS/390 V2R6 (or later) as a base component of OS/390.

A solution that combines consolidation of EE traffic at remote branches (using IBM's Communications Server products or EE-capable routers) and S/390 EE support eliminates many shortcomings of DLSw TCP encapsulation:

- EE provides data path information to network routers and virtually eliminates router overhead for SNA sessions.
- EE provides an IP connection into S/390 for encapsulated SNA traffic. Gigabit Ethernet connectivity, supported only for IP traffic, is a high-bandwidth option.
- EE eliminates termination and acknowledgement workload, thereby enabling routers to handle a much larger number of users
- EE leverages the inherent availability features of IP to provide failure-resistant SNA application access. Because EE does not terminate the session flow, it can support the IP reroute capability, maintain the connection, and switch to an alternate path without session disruption.
- EE provides traffic priority by mapping the priority of SNA packets to UDP port numbers that routers within the IP network can easily use to properly handle the traffic.

Older routers may not be capable of supporting DLSw, or they may require software upgrade.

PNS replaces SLSS

The IBM System Library Subscription (SLSS) has been replaced by an online notification application, the Public Notification System (PNS). All active SLSS subscriptions were migrated to the new system, which is accessed from the Internet at <http://www.ibm.com/shop/publications/pns/>

Dynamic VIPA? No problem – with a little planning

by Jay Aiken

Jay Aiken is a Senior Software Engineer who has been with IBM for more than 25 years. He is currently a member of the SecureWay Communications Server for OS/390 Strategy and Design Department, specializing in Parallel Sysplex exploitation. Jay has 9 invention patents granted or pending, including one for Dynamic VIPA Takeover.

So you've decided to deploy your new and fancy, bet-your-business killer application on your S/390 Parallel Sysplex for unprecedented scalability and availability. The sysplex is laid out with external links physically separated, and different power lines going to different CECs. You are aware of Virtual IP addresses, available since OS/390 V2R5 (actually also in TCP/IP 3.2) and you've assigned a VIPA to your new application and defined it on the normal hosting stack. You've defined an Automatic Restart Manager policy so ARM will restart your application on one of several other MVS images in the sysplex in the event of a failure. So all is well?

Not quite yet. Suppose you lose the router attached to the interfaces on the stack hosting your application, or Joe's Plumbing and Backhoe Services gets overly ambitious next door and cuts the power to the CEC hosting your application - with shocking consequences to Joe, of course - and takes down your application? What then? Well, as you know from reading *OS/390 eNetwork Communications Server IP Configuration*, VIPAs can be moved from one stack to another,

via VARY obey executed by the system operator or by automation. However, ARM will decide at the time of failure which of several possible targets is the most appropriate. But the application and clients will not be in touch with each other unless the VIPA is moved to the stack on the MVS image where the application gets restarted. In fact, if the application binds specifically to the VIPA, it may not start successfully unless the VIPA is already there. Sounds like a problem...

OS/390 V2R8 (and higher) to the rescue! TCP/IP provides a new capability called Dynamic VIPA (also known as VIPA Takeover) that will handle this situation. The essence of Dynamic VIPA is as follows: If the application issues a bind to a specific address, and that address is not active on this TCP/IP or any other stack in the Parallel

Sysplex, TCP/IP will activate the address as a Dynamic VIPA and allow the BIND to succeed. Naturally, it would not be good to have unconstrained activation of IP addresses in the sysplex, so a little stack configuration is required, in the form of a VIPARANGE statement. The VIPARANGE statement

identifies a range of IP addresses which may at some future time be activated as Dynamic VIPAs, in effect giving permission for a future program to activate one of them.

In order to get this to work properly for your particular application and the Parallel Sysplex environment, some planning is required:

- Determine a range of IP addresses which you would like to use as Dynamic VIPAs for your Parallel Sysplex.
- Determine where in the sysplex your application will normally reside, and where it could be restarted in the event of a failure. (Define appropriate ARM policies to make this happen.)
- Make sure that you have dynamic routing for all TCP/IPs

that might host your application, so router can learn of Dynamic VIPA movement. Add an appropriate VIPARANGE statement to the TCP/IP profile for the stack that normally hosts your application, and

for all other stacks that might host it according to the ARM restart policy.

- Add a definition to your Domain Name Server linking the application name to the Dynamic VIPA.
- Start your application and let the productive work begin...



Paul Rowntree

No problem now? Well, maybe... Suppose your application does not bind to a specific IP address, but rather accepts connection requests to its specific port on any IP address? All of the above assumed that the application would tell TCP/IP which Dynamic VIPA it wanted activated by issuing a bind() function call with the right address. With OS/390 V2R8 (and higher), no problem! ARM does not restart individual job steps, but an entire package. A utility called EZBXFDVP can be configured to activate a Dynamic VIPA for you. Simply add an invocation of EZBXFDVP to your JCL package before your killer application job step, and the Dynamic VIPA is activated as part of starting the job - or restarting the job in the case of ARM.

We've barely scratched the surface here. So if this sounds good to you, take a look at Chapter 3 of *OS/390 SecureWay Communications Server IP Configuration* for lots more information on VIPAs and Dynamic VIPAs, including an entire class of applications I did not even mention. Once you do, please give us feedback on how well we explained the concepts. Parallel Sysplex is incredibly powerful, but it is also complex. To help you get the most out of your specific configuration, we want you to thoroughly understand it.

We Speak Your Language!

But what are we saying?

- **Consistent condition handling**
 - Interlanguage communication
 - Year 2000 support
- Performance enhancements
 - **Increased efficiency & portability**
 - Centralized diagnostics
 - Ease of implementing client/server applications

We're saying LANGUAGE ENVIRONMENT

A single run-time environment for high level languages

So.....whether you speak

PL/I

COBOL

C/C++

FORTRAN

We Know What You're Saying!

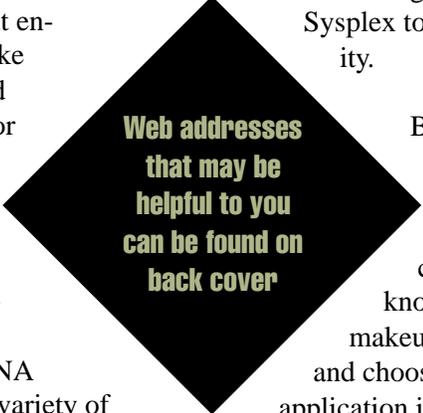
The previous issue of Hot Topics (GA22-7431-00) has lots more info about Language Environment. See page 31.

TCP/IP application workload balancing in the S/390 Parallel Sysplex

by Mac Devine

Mac Devine has worked on a variety of APPN/HPR functions and was the Chief Programmer of several major releases. He is currently a member of the Communications Server for OS/390 Strategy and Design group. He served as the Chief Designer for Communications Server for OS/390 V2R7.

For years, application workload balancing in a Parallel Sysplex environment was considered a function of SNA's Generic Resources, but recent enhancements now make application workload balancing a reality for TCP/IP applications as well. While the task of workload balancing in a TCP/IP environment is not nearly as straightforward as SNA because of the wide variety of TCP/IP applications, the value of workload balancing is consistent across both application protocols. By distributing workload among replicated applications in a Parallel Sysplex, you can make efficient use of network resources while significantly enhancing the overall availability of the data center. In order to fully appreciate these benefits, two concepts, as well as their relationship with one another, need to be understood.



Horizontal growth and single system image

On the surface, you might think that the only goal of application workload balancing in a Parallel Sysplex is simply to “spread the work” among replicated applications, but the true goal is to spread that workload in a manner which factors in the health of individual application instances as well as to allow additional application instances to be added non-disruptively to the Parallel Sysplex. In other words, you don't want to distribute workload to application instances that are not available or performing poorly, and you certainly don't want to bring down your Parallel Sysplex to add needed capacity.

By making the S/390 Parallel Sysplex look like a single system to the client, the burden of knowing the physical makeup of the data center and choosing the appropriate application instance shifts from the client to the Parallel Sysplex itself. The Parallel Sysplex provides a single system image by having either a single name, which represents the replicated application instances (DNS method) or by a single IP address, which represents the replicated application instances (cluster method). By allowing applications to be moved or added to new locations within the sysplex and still be considered part of the same single system image application group, the data center can grow without impact to the other sysplex members, the WAN or the clients. The XCF Dynamics function, delivered in OS/390 V2R7, allows for this horizontal growth of the data

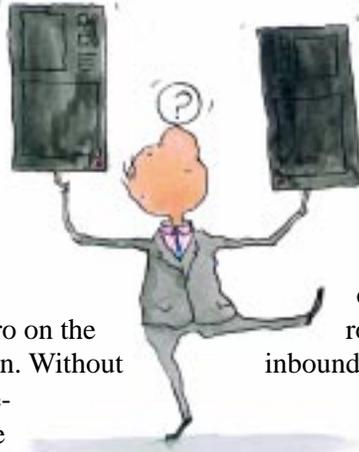
center. XCF dynamics introduced the concept of a single local IP address for connectivity to all other stacks in the Parallel Sysplex. In other words, XCF IP links under XCF dynamics look very much like a shared-medium LAN such as Ethernet or Token Ring, but all point-to-point connectivity between stacks are dynamically defined via the MVS XCF facility, which is keeping track of the members of a particular XCF group and allowing the members to communicate with one another.

DNS method

A DNS method for TCP/IP application workload distribution uses a single name to represent the replicated application instances within the Parallel Sysplex. There are two types of DNS methods available in the marketplace today: DNS weighted round-robin, available on a variety of platforms, and DNS/WLM support provided by S/390. The DNS/WLM support available on OS/390 V2R5 (and above) uses input from the MVS workload manager (WLM) to factor in the current health of the application instance in terms of CPU consumption and availability (that is, is it currently up and running on a particular stack?). This provides a competitive advantage over the round-robin approach, which does not factor in these variables.

Because of the time needed to perform a DNS resolution for every connect request from the client in the DNS/WLM model, we generally recommend it only for long duration connections like Telnet and FTP.

Also consider whether clients and/or the primary DNS (that is, a DNS which is more local to the client than the S/390 DNS) cache the IP address even though the S/390 DNS supplied a time-left value of zero on the initial DNS resolution. Without the honor of the time-left value of zero, the availability of the Parallel Sysplex is severely compromised, since the client may attempt to reconnect using the IP address of a “down” application server. The VIPA takeover function in OS/390 V2R8 significantly improves the availability in this scenario since the cached IP address of the server can be a dynamic VIPA address, which can automatically be recovered on a backup stack. However, the WLM information is not factored into the reconnect decision.



Paul Rowntree

Cluster method

A cluster method for TCP/IP application workload distribution uses a single IP address to represent the replicated application instances.

A variety of implementations are available in the marketplace, including IBM’s Interactive Network Dispatcher (IND) and Cisco’s Multinode load balancing (MNLB). Both implementations use an IP address advertised as being owned by a particular router and therefore all inbound connection requests for that IP address are directed to that particular router. On arrival at the distributing router, the connection is forwarded to the destination server chosen, based on MVS WLM information obtained via feedback agents that are running on the S/390. The destination server processes the inbound connection request, since the S/390 TCP/IP stacks have also been defined with a loop back alias address which corresponds to the cluster address owned by the router. Cisco’s Multinode load balancing has the advantage of being able to physically reside anywhere in the IP network (the IBM Interactive Network Dispatcher must be on the same IP subnet as the S/390 servers in the Parallel Sysplex) as long as there are only Cisco routers up-

stream to the datacenter. This is because the Cisco routers communicate via their Dynamic Feedback Protocol (DFP) to enable forwarding of the data packets to the “real” destination server instead of sending all inbound data packets to the router owning the cluster address. It is our intention, in a future release of OS/390, to have a function called the Sysplex Distributor also use the cluster method to allow for sysplex-wide VIPAs corresponding to individual ports, thus allowing application-centric workload distribution and providing workload distribution with total independence from network connectivity. The Sysplex Distributor may also be able to factor in Quality of Service and Policy information into the workload distribution decision in addition to WLM information.

Improved TCP/IP performance metrics

by Richard Tobacco and the OS/390 Performance Team

Members of the SecureWay Communications Server for OS/390 Performance Team are Dan Patel, Bob Perrone, David Brooks, and Bala Rajaraman.

The combined metrics of system throughput, response time, and CPU utilization demonstrate TCP/IP improvements since V3R2. This performance data was collected in dedicated, controlled system environments. Primitive workloads were used to track communication stack performance trends for representative real-world workloads. Three primitive workloads were used.

The Request-Response primitive correlates to interactive transactions. There has been up to a 16X improvement in transaction rates between the old V3R2 stack and recent releases of the new TCP/IP stack.

Connection Request-Response is the primitive that represents establishing a connection prior to each data transaction. This too shows a 16X improvement in stack performance when measured on a 10-way G4 processor.

The Workload Stream primitive is representative of file transfers. Between the old TCP/IP stack and the newer V2R6, V2R7, and V2R8 releases there is up to a 6X improvement.

TCP/IP Primitive Operation	V3R2	OS/390 V2R5	OS/390 V2R6 - 8
Request - Response (Transactions/Sec)	1674	19267	25647
Connection Request - Response (Transactions/Sec)	320	3110	5179
Workload Stream (Megabytes/Sec)	28	80	177

Note: Measurements taken on a 10-way G4 processor.

Performance of CS for OS/390 in any environment can be improved significantly by proper tuning. Three informational APARs highlight important tuning parameters and describe recommended parameter values and techniques that can maximize performance. These three informational APARs are: II11710, II11711, and II11712

To access the INFO APARs from the Web:

<http://www.software.ibm.com/network/commsserver/support/apars/csos390ip.html>

At the Web site, select 'TCPIP' as the protocol, select 'All Types' for APAR Type: find and select the APAR, click on 'Get more information' to retrieve the APAR text.

Performance tuning tips and capacity planning information for CS OS/390 is also described in our recent SHARE presentations, available at:

<http://www.share.org/proceedings/sh93/shares99.html>

Click 'Sessions' button on left. Click number 3 at top. Select the Session number to download. The CS/390 Performance Sessions from 8/99 SHARE are:

- 3915 TCP/IP for MVS Performance Tuning Tips and Capacity Planning
- 3631 TCP/IP and SNA Performance: Present and Future
- 3819 Gigabit Ethernet Performance and Tuning for the Mainframe

ATTENTION: Language Environment Run-Time Options

The ABTERMENC run-time option has a new default sub-option. It's been changed from RETCODE to ABEND. So if you're migrating to Release 9 and want to retain RETCODE, you'll have to code it. But we've always recommended the ABEND sub-option and if you've selected this in the past, relax. You have one less thing to do when migrating. We've done it for you.

A roundup of IP security

by Lin Overby

Lin Overby is a Senior Software Engineer at IBM Corporation. He works in the Strategy and Design area of SecureWay Communications Server for OS/390. Lin is the lead designer for new offerings in the area of TCP/IP security which include IPSec/ Internet Key Exchange and TN3270 SSL.

The SecureWay Communications Server for OS/390, along with other elements of OS/390, provides numerous security technologies to protect your mission-critical data in a TCP/IP environment. This article is about these S/390 technologies and how you can use them to securely participate in the ever-expanding world of e-business.

Cryptography -The foundation of good security

Good security methods begin with cryptography. Encryption protects sensitive data from being read by anyone other than the intended receiver. Cryptographic authentication and data integrity services allow communicating hosts to detect if data is altered in transit. Public-key cryptography can identify and authenticate hosts or users. Once a secure session is created, successful data authentication and decryption occurs only if both hosts have the correct session keys.

End-to-end security

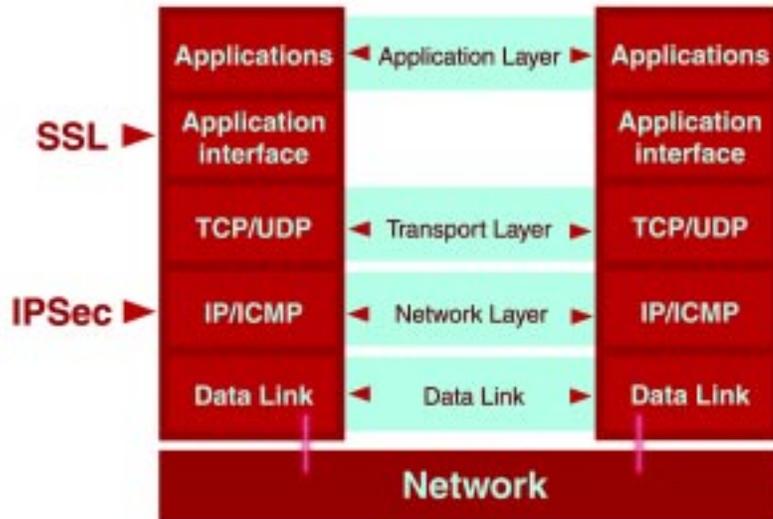
Cryptographic security solutions can be applied to a portion of the data path or end-to-end, whichever is appropriate for your security policy. Generally, the greatest degree of security is provided when cryptographic methods are used end-to-end. However, if an enterprise considers only portions of the data path untrusted (such as the Internet) it may be adequate to protect only the untrusted segment with cryptography. OS/390 offers security protocols that can be configured to protect portions of the data path or the entire data path.



Paul Rowntree

Workload-based security deployment

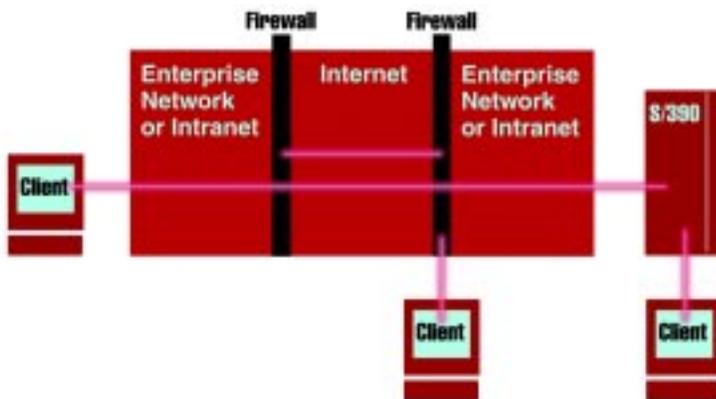
In making a security protocol selection, an important consideration is the application workload to be protected. In order to illustrate this concept, it is helpful to understand where various protocols are implemented from a layering perspective.



The network layer is the lowest layer in the protocol stack where end-to-end security over multiple hops can be applied. Network layer security protocols provide blanket protection for upper-layer application data without requiring modification to the application. IPSec is implemented at the network layer and can be used to provide authentication, integrity, and data privacy between any two IP entities. IPSec can protect a segment of the data path (for example, between two routers), or it can secure the data path end-to-end. Management of cryptographic keys and security associations can be manual or automated via an IETF-defined key management protocol called the Internet Key Exchange (IKE).

IPSec can protect selected traffic or all traffic. Selectivity can be based on various combinations of source and destination IP address, port, and protocol. This selectivity can be used to avoid the overhead of multiple security protocols when alternate security protocols are used to secure specific applications. For example, you might want to exclude Web traffic (based on well-known port of Web server) from IPSec coverage because you would like to use SSL.

IPSec allows the creation of Virtual Private Networks (VPNs). A VPN enables an enterprise to extend its network across a public network such as the Internet through a secure tunnel (or security association). The figure below shows some of the typical IPSec configurations. IPSec security associations are shown 1) between two firewalls, 2) between client and firewall, and 3) between client and S/390 server.



Communications Server for OS/390 IPSec supports the latest RFCs including Triple DES for strong encryption. OS/390 V2R8 support of IKE reduces the manual effort in managing and distributing IPSec session keys and also supports a non-disruptive session key refresh capability. Ease of key management is a critical factor as the numbers of IPSec hosts increase.

Secure Sockets Layer (SSL) is another popular security protocol implemented above the transport layer at the application interface layer. SSL, originally used to secure traffic between a Web browser and Web server, can also be used to secure other applications.

One such user of SSL is the TN3270 server provided by Communications Server for OS/390. Serving as a gateway between the IP network and the SNA network, the SSL-enabled TN3270 server protects the data path in the IP network from the TN3270 client all the way to the S/390 TN3270 server. If the TN3270 server resides on a different host from the target SNA application, SNA Session Level Encryption can be used to secure the SNA portion of the data path. In OS/390 V2R8, the TN3270 SSL support was augmented to support client authentication and pre-login access control using RACF digital certificate support. Before users can receive a logon screen from TN3270, they must provide a certificate from a trusted certificate authority that is authenticated as part of the initial SSL session setup. Additionally, this support also verifies that the user associated with the certificate is authorized to access the TN3270 server port.

For new applications, security can be built-in. One method of building security into the application on S/390 is to use OS/390 System SSL. Newer versions of network services such as SNMPv3 and Dynamic DNS have security built into the protocol using standards-based specifications for secure interoperability.

Has it been ages...

since you found a hit
in the OS/390 Collection Kit?
Now in just ten pages
and with one quick read

our **NEW**
"OS/390
Collection
Summary"

alphabetically
expeditiously
helps you
find the
bookshelf
you
need

See
page
30

Is OS/390 TCP/IP ready for DB2?

by Hugh Smith

Hugh Smith is a member of the DB2 development team specializing in DB2 Client/Server Performance Analysis.

Introduction

The answer to the question posed by this article's title is: **YES!**

Since the general availability of DB2 for OS/390 V5, the required TCP/IP support within OS/390 has significantly improved, especially in the area of performance. Our goal has been that the TCP/IP within the OS/390 operating system would either meet or surpass the performance of SNA so that your decision on which protocol to use could be based on factors other than performance. The TCP/IP support provided with OS/390 V2R6 has met this goal, especially when using it to access DB2 via Distributed Relational Database Architecture™ (DRDA) requesters such as DB2 Connect V5 or V6.

While most of the earlier improvements occurred in the overall TCP/IP protocol stack within OS/390, most of the recent improvement in the OS/390 TCP/IP support has occurred in the network device driver support. This improved device driver support is called Multi-Path Channel + (MPC+). This device support was intended to reduce CPU consumption compared

to other device drivers, but it can also be used to support any network device that is channel-attached to OS/390 for both TCP/IP network flows as well as SNA flows. A reduction in CPU consumption realizes a gain in network throughput.

Before reviewing the study, which was performed to verify the previous claim, it is important to review some tuning tips and recommendations that will optimize the use of those network devices. Since the topic of tuning distributed, client/server, network-computing type applications can be somewhat lengthy, you might want to check out a document at <http://www.software.ibm.com/data/db2/performance/tcpsna.htm>, which covers this topic in more detail. The focus of tuning recommendations in this article will be on how to best utilize the network protocol and devices.

TCP/IP is a very simple protocol and because of this, there are few tuning recommendations. It is also a very flexible protocol. One of the most important factors in the network is understanding the flow of one's requests through the network. Most, if not all, TCP/IP networks are routed networks. That means your messages with the DB2 server could traverse quite a few nodes along the way and back. Understanding your network path and trying to optimize it as best as possible is an important step to client/server application tuning. The fewer the nodes the better, because each node will add a latency to the overall network flow. On all systems but WIN95/98, a command to trace the route to a host is available, such as traceroute (UNIX), tracert (NT), and tracerte (TSO),

otracerte (Unix System Services). This command gives the user a set of times to traverse the network through each node as well as all the nodes that you must traverse to get to your host. The last line of the output is the total send/receive network time for each message. Thus, if your application is very chatty (lots of SQL flows), you will have to take the number of messages for an application transaction times the average of times on the last line of the trace route command and you should get a basic, small message, network application time. For instance, if the trace route command said 30 ms (milliseconds) and you have to exchange 100 SQL statements, the network cost alone is about 3 seconds. That doesn't even take into account the amount of time in DB2 to process the requests. Now the operative word here is "optimize" your path to the DB2 server. Not everyone can have the best path to DB2. That will require a delicate balancing act on the part of users, systems personnel, network personnel, and application developers.

The next area to investigate for optimizing network traffic is the IP Maximum Transmission Unit (MTU). The setting is done in OS/390 via the GATEWAY routing statement. Never, ever, set the MTU size to DEFAULTSIZE. This will mean that all messages in and out on the link it is specified for will be at most only 576 bytes. Thus, if DB2 wants to send a 32K cursor block back to the client, over 60 messages will have to be generated. Thus, for Token Ring and FDDI, sizes of 4400 and 4000 are recommended,

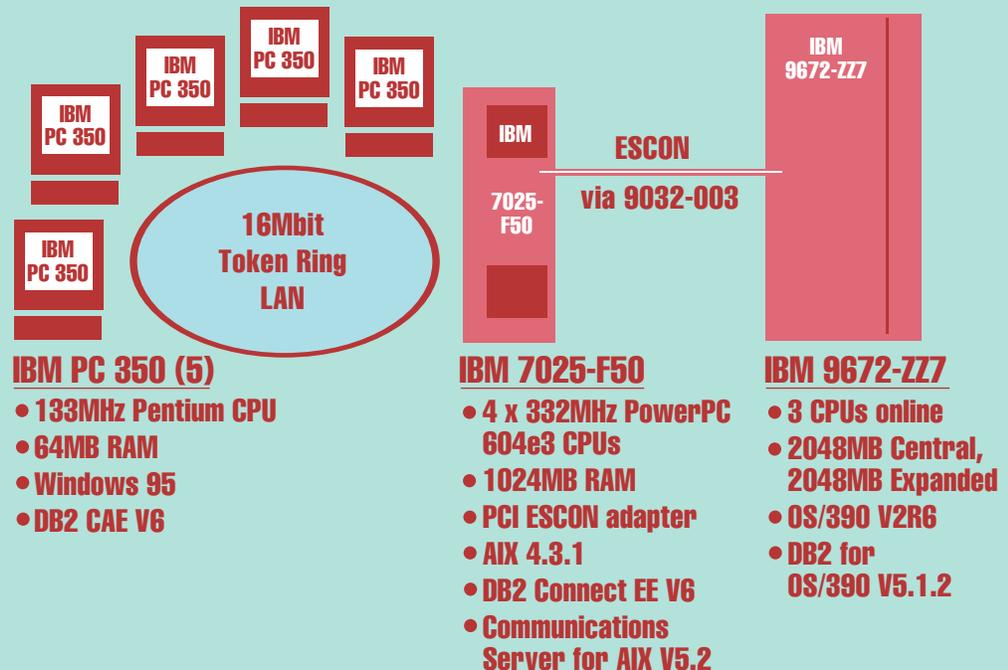
respectively. For Ethernet, no matter what speed, do not set anything other than 1500. On WIN32 systems, the TCP/IP protocol stack constantly negotiates this size, which is a terrible overhead. To turn off MTU renegotiation, you will have to go into the registry editor (more on the TCP/IP registry setting later). As to ESCON adapter, the recommended settings go as high as 4096 for the CLAW protocol (more on this later) and can go up to 61440 for the MPC+ device drivers.

Finally, in order to perform flow control across the network, TCP will send up to a specified buffer size and pause until it receives an acknowledgment from its partner for any byte within the block. The “window” is then slid forward to the acknowledged byte, and sending is continued until this buffer is filled again. The same is true on the receiving end of a connection. There is a receive buffer. These buffer are called the TCP Send/Receive Spaces. It has been found in studies, that for DB2 Connect and its clients (or going direct to DB2 on OS/390), that a send/receive space size of 65535 is the best. This is altered in the OS/390 TCP/IP profile dataset on the TCPCONFIG stanza. For AIX systems, a command called “network options” or “no” is the way to do it. You will have to consult your UNIX systems documentation on the way this can be done on your UNIX system. For WIN32, systems this requires a registry edit. For the details on all the possible TCP/IP registry settings, you will have to consult the Microsoft Technical Support article Q120642 for NT and Q158474 for WIN95/98 (the default is only 8192). The setting of your TCP Send/Receive Spaces sizes does

require some analysis of its impact. The larger the space sizes are in a congested network, the greater chance data will not be acknowledged soon enough causing data to be resent.

Configurations

The following image gives a pictorial view of the environment used to perform the study:



On each client workstation, 4 simulated virtual users ran a workload application for a total client population of 20 users. The workload is described in a document which can be found via: http://www.software.ibm.com/data/db2/performance/tcpsna05.htm#Section_4.0.

The workload used during this study differs from that description in the following ways:

- The SQL processed against DB2 came from the workload application program on the client workstation rather than from stored procedures.
- The workload application program was written in C and used the DB2-CLI application interface. The program did not do any windowing.

The kinds of transactions and their percentages of the workload mix were the same.

Measurement Process

For the purposes of this study, the workload application, the number of virtual users per client workstation, the DB2 client workstations, the connectivity (TCP/IP) used between the client workstation, the DB2 Connect EE V6 AIX server, the physical connection with the mainframe, and DB2 for OS/390 V5.1.2 server and the application workload database remained constant. The factor that was varied during the study was the protocol used between DB2 Connect EE V6 on AIX and DB2 for OS/390 V5.1.2 on the OS/390 V2R6 system. The protocols studied were:

1. SNA via eNetwork Communications Server for AIX V5.2 and Advanced Peer-to-Peer Networking (APPN) enabled, no High Performance Routing (HPR) enabled, Multi-Path Channel (MPC) protocol over an ESCON connection (HPR is not supported on ESCON connections to the OS/390 system in the current eNetwork Communications Server for AIX V5.2 product)
2. TCP/IP via the Common Link Access to Workstation (CLAW) channel protocol over the same physical ESCON connection
3. TCP/IP via the MPC+ channel protocol over the same physical ESCON connection

The actual physical link, an ESCON connection with the mainframe via the PCI ESCON adapter, was also constant throughout the study. The only difference between the protocols studied was either the device driver software on both the AIX and OS/390 systems and/or the messaging protocol, SNA/APPN or TCP/IP sockets. This connectivity configuration was also chosen for the study because it is very similar to that used by many channel-attached routers, although DB2 Connect software would not be running on such devices. The complete details on how the ESCON link was defined can be found at the following URL:

<http://www.software.ibm.com/data/db2/performance/xtcpsna.htm> .

For each of the connection types studied, two measurement runs were performed where the 20 virtual users executed the workload application program for a 10 minute duration. After 30 seconds into the measurement interval, RMF™ was started to record the CPU utilization

on OS/390 and vmstat was started to record the CPU utilization on AIX for a duration of 9 minutes. At no time during the measurement runs did the CPU utilization on the workload client workstations ever reach 100%.

Results

The focus of this study was to verify the claim that TCP/IP on OS/390 is at least equivalent to SNA. Thus the OS/390 CPU utilization measured during the two different intervals for each connection type was used. Even though the CPU on the AIX workstation varied for each of the connection types, once the request message was received into OS/390, it was the overall CPU cost of processing the request on OS/390 that was of interest. The OS/390 CPU utilization did include the CPU required by DB2 to process the request. However, since the SQL originated from the workload application program on the client workstations, a CPU communications cost occurred per SQL request. In a measurement run independent of this study, the DB2 accounting trace was used to determine the average number of messages per transaction, which was 24.45 messages. Also, the average transaction sent 7586 bytes into DB2 on OS/390, which resulted in DB2 on OS/390 returning 6803 bytes.

The following table documents the study's measurement results. The External Throughput Rate (ETR) is the total number of completed successful transactions by all the

virtual user clients divided by 600 seconds (10 minutes). The AIX CPU utilization is being provided in the report because it was the limiting factor for the base SNA case. The OS/390 CPU utilization is used to derive the Internal Throughput Rate (ITR) which is the measured transaction throughput rate normalized by the amount of CPU consumed under OS/390. The ITR is calculated by dividing the ETR by the measured OS/390 CPU utilization, e.g. $41.33 / 0.1421 (14.21 \%) = 290.86$ transactions per second. The ITR delta is then the percentage difference between the TCP/IP protocol used and base SNA measurement. Only the results from the measurement run which yielded the best ITR value for each connection type was used.

Conclusions

This study does indeed verify the claim that TCP/IP under OS/390 is ready for DB2. Even though the TCP/IP via CLAW connection was degraded by 1.6%, the degradation in ITR is so small one could almost claim the TCP/IP CPU cost is equivalent to SNA. However, the MPC+ device driver used by TCP/IP is the clear winner. Not only is DB2's performance better with TCP/IP, TCP/IP is the focus for future performance enhancements. Those enhancements include OSA Express device driver support for Gigabit Ethernet, which uses new Queued Direct Input Output (QDIO) architecture.

Disclaimer

The information contained in this article has not been submitted to any formal IBM test and is distributed AS IS. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any performance data contained in this technical report was determined in a controlled environment, and therefore, the results that may be obtained in other operating environments may vary significantly. Users of this information should verify the applicable data for their specific environment

Protocol	External Throughput Rate (ETR)	AIX CPU Utilization (%)	OS/390 CPU Utilization (%)	Internal Throughput Rate (ITR)	ITR Delta (%)
SNA only	20.79*	97	8.42	246.89	base
TCP/IP via CLAW	29.79	24	12.26	242.96	-1.6
TCP/IP via MPC+	41.33	34	14.21	290.86	+17.81

*ETR for SNA was limited due to the CPU capacity of the AIX platform.

To learn about viewing,
printing and ordering OS/390
information, check out
[http://www.ibm.com/s390/
softcopy](http://www.ibm.com/s390/softcopy)

Introducing LOOKAT, a prototype message help facility for OS/390

by Jim Steipp

LOOKAT this ...

Frustration, frustration, frustration ...
Something is happening to your
system and you're trying to decipher
a cryptic OS/390 message and you
need an explanation FAST, but
where is it? You don't know which
book to look in. Or you don't have
the right book. Or the right book is
at the wrong level. More frustration
... but LOOKAT can help.

Message explanations at your fingertips

Currently under test by IBM,
LOOKAT is a prototype of an
OS/390 help system that gives you
full message explanations when you
need them, online, always at your
fingertips. Two components make
up the prototype. The first is a
message data base created from the
messages books available on the
OS/390 Online Library Collection
(SK2T-6700) on CD-ROM. The
data base is simply a BookManager®
repository set up on an OS/390 host.
The second component is the code
that locates the correct message and
displays it on your screen.

Using LOOKAT

Using the simple command
LOOKAT, followed by the message
id, you can retrieve the explanation
(and associated information) of any
OS/390 message that you might
receive. LOOKAT uses
BookManager READ/MVS and the
OS/390 messages and codes books
to locate and display the
explanation. Once the message is
displayed, you can use all of
BookManager's power to find any
additional OS/390 product
information that you may need.

You can enter the LOOKAT com-
mand in a variety of places:

- In TSO at the "READY" prompt
- From the ISPF Command panel
(Option 6 From any ISPF command
line (prefixed by the letters "TSO"
and a blank space)

If you are running an ISPF applica-
tion, you can also split the screen
and have your message explanations
appear just there.

How you obtain LOOKAT

All the items you need to install and
set up LOOKAT reside on one of the
discs in the *OS/390 Online Library
Collection*. These items are:

- **LOOKAT.EXE** - Is the code that
handles the LOOKAT command and
finds the correct message explana-
tion.
- **LOOKAT.DIR** - Is a list of the
messages books that **LOOKAT.EXE**
uses to locate the messages and
codes books.
- **LOOKAT.IND** - Is a large table that
the **LOOKAT.EXE** uses together
with **LOOKAT.DIR** to locate a
specific message explanation.
- **EOXVSTRT.EXE** - Is a
BookManager invocation EXEC.
- **LOOKAT.ME** - Is an online, ASCII
version of the instructions for using
LOOKAT.

We need your help

We would like you to help us by
looking at the LOOKAT prototype
and telling us what you think and
how we could improve it.
If you would like to help, please
contact the LOOKAT development
team at: jfsteipp@us.ibm.com

We will give you specific instruc-
tions on how to obtain and use the
prototype and on how you can
enable the BookManager versions of
your messages books for LOOKAT.

The growing wizardry of OS/390

by John Sefcik

If you'd like some help with some of those complex tasks that you do infrequently, then check out our Internet-based wizards. We use the term "wizard" to refer to them generically, but they actually have the word Assistant in the title, and it is often preceded by a word like Planning or Configuration. For example, there is the OS/390 Installation Planning Assistant, and the S/390 Parallel Sysplex Configuration Assistant.

The wizards provide a simplified approach to completing a set of OS/390 tasks, such as planning, installation, and configuration. The wizards provide customized information based on your individual requirements and answers. The wizard asks you for information, and then takes you to the next step. The wizards may also simplify your planning and configuration needs by exploiting recommended values, and by providing customized checklists that reduce the number of steps and number of information sources you need to refer to. There is help in the wizards if you need some background information. And, if you want to read more in-depth background material, there are often links to books that provide additional, related information.

Our current wizards are Internet-based, and do not directly perform any tasks on your OS/390 system itself. Instead they generate things like tailored instructions, or jobs, policies, or parmlib members that you can upload to OS/390 and use.

The list of wizards that will be available in the OS/390 Release 9 time frame are:

OS/390 Installation Planning Assistant, which builds customized installation plans for customers who are migrating from one release of OS/390 to another.

S/390 ServerPac Ordering Assistant, which assists with OS/390 ordering by building an order request for selected products in the ServerPac catalog and submitting it to IBM for processing.

S/390 Parallel Sysplex Configuration Assistant, which builds tailored instructions, jobs, and parmlib members for configuring a sysplex. It is intended for customers who are setting up a Parallel Sysplex for the first time.

OS/390 Planning Assistant for e-business, which recommends a set of products and features for S/390 e-business based on the user's objectives and current configuration. It is intended for technical managers responsible for e-business planning.

OS/390 UNIX Configuration Roadmap, which lists the actions required to set up OS/390 UNIX System Services. For assistance with Step 1, basic UNIX setup, we have provided a wizard. Over time, we will add new wizards to support more tasks in the roadmap, such as file system set-up and initial security services.

To try out these wizards, go to <http://www.ibm.com/s390/os390/bkserv/wizards.html> and click away. As we develop new wizards, or enhance these existing ones, you can always find the latest information and copies by checking out this wizards web page. We are looking to continue to expand our efforts into areas that will help customers deal more easily with a wide range of more complex system tasks or infrequently performed, but time consuming, system tasks.



Jodi Holmes

Easing your way through those large OS/390 collection kits

by Shirley Swenson

Have you ever expressed the opinion that looking for a needle in a haystack might just be easier than locating a book in our humongous Online Library collections? But did you know that our Online Library collections include several retrievability aids?

For example, look at the printed booklet packaged with the collection. Yes, hardcopy! But in this case, it's helpful because you can use it as an index to the library. The "Table of Contents" lists all bookshelves alphabetically and includes the number of the disc where they reside. And if you know a book's order number, you can find out where it is by looking at the section called "List of Books in This Collection - Arranged by Order Number."

And if you'd rather not deal with hardcopy at all, you'll be pleased to know that the collections include an online version of the booklet—on each disc, in its own bookshelf. Open the bookshelf called "Booklet for OS/390 Collection Bookshelf" (filename EZ239Syq, where y is the year and q is the quarter). You will see the booklet called "OS/390 Collection Booklet" (filename EZ239Byq) Open the booklet and you can search for any keyword that might appear in a book title or product name. Be aware that you are searching only the collection booklet, but you will quickly see which books contain the search words in their titles and on which discs they reside.

An ASCII (text only) version of the booklet named README.ENG and most recently a PDF file named README.PDF are also available. They are in the root directory on each CD-ROM. You can display the ASCII version as is, or use your favorite word processor to search it as plain text. You can print your own high quality hardcopy of the booklet from the PDF file.

We also include an HTML version of the booklet on the CDs. This version, which is also in the root directory, contains "active links" to the books. If you have the disc containing the book you want in your CD drive, you can open a book in the list just by clicking on it.

In December 1999, we began distributing a hardcopy summary of the product libraries on the collection to help you locate the library you want. This booklet is called

"OS/390 Collection Summary".

Online versions of this summary are included on the Booklet for OS/390 Collection Bookshelf, and in text and PDF format in the root directory on each CD (named SUMMARY.ENG and SUMMARY.PDF).

You can also do a full content search of the books on any disc. Go to the bookshelf called "All-Disc x Bookshelf" (where x is the number of the disc). The search will take place across the contents of all the books on that disc.



Kimberlee Mule

Saving your 'BookManager Library Reader for Windows' sort order

By default, BookManager Library Reader displays bookshelf and books by those obscure filenames, not by descriptive titles. When you changed the order to sort by titles of bookshelves and books, you couldn't save the change for future sessions. Now you can. And here's how.

With the 2.0.1 or later version of the BookManager Library Reader for Windows installed, after you select the type of sort you want from the "Sort by" selection of the View menu, go to the Options menu and select Save Window. Your selected sort, for instance by Title, is the default the next time you open the Windows Library Reader.

Here's where you can find Hot Topics

- **Hardcopy of the Newsletter is included in each "OS/390 Installation Planning Kit" that you get free of charge when a new release of OS/390 is made generally available (twice a year).**
- **We will continue to distribute hardcopy at user group meetings like SHARE.**
- **Softcopy is available on the Web in PDF format (readable through Adobe Reader) at <http://www.s390.ibm.com/ftp/books/os390/pdf/hottopic.pdf>.**
- **It is available from the OS/390 home page (<http://www.s390.ibm.com/os390/>), under publications in BookManager format. This version will be cumulative, meaning that all past issues will also be available at this location.**
- **This softcopy cumulative edition is also available on each copy of *OS/390 Online Library Collection (SK2T-6700)*, Disc 1.**
- **It is also on each copy of *OS/390 PDF Library Collection (SK2T-6718)*.**
- **It is also available to order hardcopy through your regular mechanisms at a nominal fee.**

The URLy bird gets the info

by Kevin Godbey

Because you've often told us that the sheer quantity of information on the Web sometimes prevents you from easily finding what you need, we provide you, with each new issue of the *Hot Topics Newsletter*, a list of web addresses we think will be helpful to you.

Enterprise Server essentials!

A non-technical guide for executives whose business goals cannot be compromised by I/T-imposed limitations (A GREAT analysis of S/390® vs. UNIX platform capabilities).

<http://s390.ibm.com/marketing/gf225122.html>

Are you ready for e-transaction processing?

The Design Center for e-business Transaction Processing is waiting to help YOU with those complex e-business applications.

http://www.ibm.com/servers/design_center/

Customer experiences... get a REAL feel for the REAL deal!

Read how S/390 is helping customers reach their business goals by providing faster, more reliable information processing with significant cost savings.

<http://www.s390.ibm.com/customer/>

Providing solutions

S/390 provides solutions that solve business problems, add important functions to S/390 products and simplify installation. Read success stories to learn how our customers are using IBM and solution developer products to enhance the S/390 platform to solve their business challenges.

<http://www.s390.ibm.com/solutions/>

New Technical Support pages

Take a look at comprehensive information on fixes, discussion groups, downloads, 'hints and tips', technical publications, etc.

<http://techsupport.services.ibm.com/support/s390/>

Downloads specifically for OS/390

The latest function for developing and supporting e-business applications is at your finger tips.

<http://www.s390.ibm.com/os390/downloads.html>

Essential OS/390 migration information

Read this "Planning Guide for Multisystem Customers: OS/390 Coexistence and Planning Considerations Through the Year 2000" to plan your future migrations.

<http://www.s390.ibm.com/stories/year2000/coexist.html>

S/390 and OS/390 Security – A key element of e-business

Read about hardware isolation functions, system level security, network level security, transaction level security, cryptography support and a whole lot more before designing those e-business applications.

<http://www.s390.ibm.com/security/>

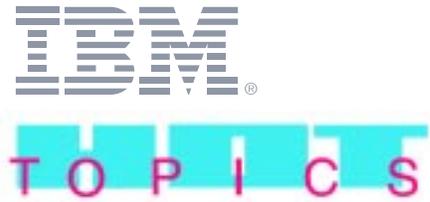
Can't find the site you need?

Go here to e-mail us...we'll help you find it!

<http://www.s390.ibm.com/os390/webqs.html>



Tracy Bondi



An OS/390 Newsletter for Release 9 February 2000

Executive Editor Bob Ward

Managing Editor Dick Wagenaar

Design Director Gene Posca

© International Business Machines Corporation, 1999, 2000

Produced in the United States of America
2-2000

All Rights Reserved

The OS/390 Hot Topics Newsletter is published twice yearly as a service to OS/390 customers, providing articles of interest to the OS/390 community.

The information in this Newsletter is presented as is and is based on the best knowledge of the authors. No warranty is provided (neither expressed nor implied).

IBM may not offer the products, features or services discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM representative for information on the products or services available in your area.

© IBM, AIX, APL2, BookManager, DB2, ESCON, Language Environment, OpenEdition, OS/2, OS/390, Parallel Sysplex, RACF, RS/6000, S/390, SecureWay, VTAM are registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

™ Distributed Relational Database Architecture, DFSMSdfp, DFSMSHsm, DFSMSdss, MVS, RMF are trademarks of International Business Machines Corporation.

Java, and all Java-based trademarks are trademarks of Sun Microsystems, Inc., in the United States, or other countries, or both.

Microsoft, Windows, Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark in the United States and other countries licensed exclusively through The Open Group.

Other company, product, and service names may be trademarks or service marks of others.



GA22-7431-01

TOPICS

OS/390 Release 9 SUPPLEMENT February 2000

UNIX on OS/390: A Powerful Blend

by Duna Williamson

Duna Williamson, Senior Information Developer at IBM®, Poughkeepsie, has been writing about MVS™ and OS/390® for 19 years. For the past five, she has concentrated on OS/390 installation, migration, and customization, trying to make these tasks easier for you. For ten years, she has been a representative to the GUIDE user conference and now SHARE. She is currently busy helping to develop an OS/390 UNIX® Configuration wizard and coordinating activities to improve the UNIX and IP libraries.

The merging of two platforms, especially two as disparate as UNIX and MVS (now called OS/390), is bound to cause confusion for those who have to maintain, debug, and write applications for the merged system. Users generally have experience in either MVS or UNIX, and neither side is quite sure what to make of the other. Terminology is different: an address space is no longer an address space but a process; a thread is no longer a thread but a task, and a started task is a daemon. A daemon? Not only how we talk *ABOUT* the platforms but how we talk *TO* them is different. One is upper case, one is lower case. One uses a shell, the other has TSO/E and ISPF. Applications are written in different languages and jobs are submitted for processing in different ways. And so forth...

But the team who manages OS/390 workloads and installs e-business products must be equally comfortable in both environments. That team needs to understand that OS/390 is a certified UNIX platform, that there is a history behind UNIX, and that, though the differences may be confusing at first, tremendous value can be derived from the system if people understand the concepts and operational

procedures of both and can translate between them.

The following pages address this need by explaining:

- What UNIX is
- Why IBM implemented UNIX on the S/390 platform and how we did it
- Some UNIX history and concepts
- What's different and what's the same between UNIX and MVS

Table of Contents

- What is a UNIX System? •pg 2
- Background of OS/390 UNIX •pg 2
- Why did IBM add UNIX to the MVS Environment? •pg 3
- IBM's Implementation of UNIX •pg 3
- Some UNIX Concepts •pg 4
- Daemons •pg 4
- The Hierarchical File System •pg 5
- OS/390 UNIX Security •pg 5
- Superuser Authority •pg 6
- The Multiple Ways to Use OS/390 UNIX •pg 6
- Comparing MVS and UNIX Concepts and Terms •pg 6
- Tables •pg 7-13

The main audience for this supplement is the MVS system programmer for whom OS/390 UNIX System Services (OS/390 UNIX) was, at times, an unwelcome addition to the computing environment. The table on page 7 is for everyone who runs OS/390. Don't miss it; it compares and contrasts tasks and concepts in UNIX and MVS in terms everyone will understand. The topics come directly from discussions that took place on the MVS-OE mailing list. This list is an excellent source of help for OS/390 UNIX. The participants are customers

other's questions; they receive all questions, answers, and comments as e-mail. To subscribe (and perhaps participate) send a note to listserv@vm.marist.edu. In the body of the note put:

```
sub mvs-oe your_first_name your_last_name
```

Once you are subscribed, you will receive further instructions on how to use the mailing list.

What is a UNIX System?

A UNIX person might respond that a UNIX system has directory and file structures that are similar to Windows or other PC systems. It has TCP/IP to connect to other systems. It has a set of system calls that perform functions that are similar to the functions other UNIXes provide. It allows users to telnet or login to a shell (the UNIX command-line interface) to interact with the system. It has certain utilities that the UNIX user expects to find in the work environment. UNIX application developers expect C or C++ compilers and X-Windows, a graphical interface.

OS/390 UNIX is a UNIX system because it has all these things as well as UNIX-branding. IBM's first UNIX, called OpenEdition® (a name you might remember), was one of the first UNIX offerings to meet the standards of POSIX (portable operating system interface for computer environments), an IEEE operating system standard.

Background of OS/390 UNIX

History loves a paradox. Even the evolving histories of OS/390 and UNIX are no exception. Starting as a batch-oriented operating system running on the world's largest mainframes, MVS evolved into an interactive system that, as technology advanced, can now run on machines the size of a PC. And UNIX, starting as an interactive operating system running on workstations, converged with MVS in OS/390 to run on the world's largest mainframes.

Describing the background of OS/390 UNIX means first looking at the history of MVS. For more than 30 years, MVS programs have run on IBM and non-IBM mainframes, and have invoked services that use data in MVS data sets, such as sequential, partitioned, or VSAM data sets. The programs have been proprietary, in that they have run ONLY on MVS. In the beginning, programs and input data were stored in 80-column punched cards. To process the data, jobs were submitted all at once as batch jobs. Scientific and business applications were developed that took advantage of the computer's phenomenal calculating speed. Some business-critical programs written in the 1960s are still running today.

The history of UNIX is very different. In the UNIX world, there are many varieties of UNIX, and each of them contains a few lines of code from the first UNIX. To summarize, AT&T created the original UNIX in 1969 and, in 1973, began to make the source code available outside of AT&T. Others ported the code to their platforms and added their own extensions to make their UNIX unique from others. These extensions made each UNIX platform different from others and introduced challenges to those porting applications from one UNIX platform to another. Since then, standards have eased the process of porting.

UNIX is an interactive operating system. It provides application developers with many small, simple tools that they join together to do complex things. UNIX does not have all the checks and balances that MVS has; it provides the tools you need and assumes you know what you are doing.

In 1994, in MVS SP4.3, IBM introduced UNIX, calling it OpenEdition. Adding UNIX to MVS required IBM to make changes to MVS components, such as RSM and contents supervisor, to make UNIX work. Although the first version of OpenEdition met some POSIX standards, it offered no entry to the network and did not include support for TCP/IP. In OS/390 R2, OpenEdition received full UNIX95 branding; OS/390 UNIX became a UNIX system according to the people who own the definition of UNIX, the Open Systems Foundation.

The UNIX that is in the current release supports all the C functions and command interfaces in the Open Systems Foundation's Single UNIX Specification. UNIX98 is the most recent standard, and IBM is implementing the UNIX98 functions that make sense on OS/390.

Why did IBM add UNIX to the MVS Environment?

IBM added UNIX to the MVS environment to compete for new work on S/390.

In addition to the characteristics of UNIX described earlier, portability is important — the openness of UNIX.

For programmers, portability of skills from one UNIX to another is important. If you work on a UNIX system, you can port most of your skills onto another UNIX platform. If you are familiar with AIX — to you AIX is UNIX — and you go to another UNIX system, there will be some differences, but, in general, you will feel right at home.

For applications, portability of source code from one UNIX platform to another is important. If applications consisted of only interfaces that meet the X/Open spec1170 specifications (also known as X/Open XPG4.2 or UNIX95 or Single UNIX Specification), a port would require no changes to the applications. But, in the real world, where applications contain some non-standard UNIX code, the applications will require some changes.

So, the two major reasons for UNIX on S/390 are portability of skills and portability of code. You can use the same skills on one platform and then another, and you can move function from one platform to another.

IBM's Implementation of UNIX

IBM did not port someone else's UNIX to OS/390. OS/390 UNIX is not an emulation, a subsystem, or an LPAR; it is not sitting on top of S/390 (as in a layering effect that suggests bad performance). IBM looked at the spec1170 specifications that say "these are what the UNIX interfaces are" and implemented those specifications directly into the OS/390 system as OpenEdition services, which are now OS/390 UNIX services. The OpenEdition MVS Shell and Utilities, with many of the Korn shell features, were first to take advantage of the new UNIX system services in OS/390.

To manage OS/390 UNIX, IBM wanted to use MVS system programming and operational skills. They wanted those who understand and manage S/390 availability, reliability, and security to continue to manage this new system in ways familiar to them. Therefore, there are MVS ways to manage the hierarchical file system, and SMP/E management of product code and service updates. However, UNIX programmers who expect UNIX interfaces will find those interfaces, unaware that, under the covers, traditional MVS things happen. For example, OS/390 UNIX services operate in MVS virtual storage areas called address spaces. Both the MVS system programmer and the UNIX application developer can be comfortable in this new environment.

There are many services in OS/390, such as MVS services and OS/390 UNIX services. All of these services run on MVS. Existing MVS applications continue to run, unaffected by the existence of OS/390 UNIX and applications that use the OS/390 UNIX services. All the applications can benefit from the reliability, availability, and security of the S/390 platform.

A common misconception is that application programmers have a choice of either running UNIX programs that use UNIX services and file systems or running MVS programs that use MVS services and MVS data sets. This choice between OS/390 UNIX and MVS does not exist. For example, if you are developing a UNIX application that needs to call

DB2[®], you do not have to pass a request to an MVS program that then extracts data from DB2 and passes data back to the UNIX program. You can call DB2 directly from UNIX. And, if you telnet into the shell, you have a standard UNIX interactive command-line interface like the one you are familiar with. There is no wall between UNIX and MVS, and there are no ‘sides’; you do not run ‘on the UNIX side’ or ‘the MVS side’. OS/390 is a powerful blend of UNIX System Services and MVS.

Some UNIX Concepts

MVS-experienced system programmers now installing, maintaining, and debugging OS/390 UNIX are faced with new responsibilities that require new vocabulary and new concepts. Perhaps they must interact with UNIX or TCP/IP computer personnel from whom they have been isolated before now. Differences between the familiar MVS world and the new UNIX world can sometimes seem cultural and philosophical, as well as technical. This section tries to bridge the gap.

The Kernel and the Shell

The UNIX operating system has two distinct pieces, the kernel and the shell:

- The kernel is part of the BCP element of OS/390; it sends instructions to the processor, schedules work, manages I/O, and tracks processes, open files, and shared memory, among other things. Other parts of the operating system or applications request the kernel’s services using assembler callable services (called syscalls). No work gets done in OS/390 UNIX without involving the kernel.
- The shell, the interactive interface with the kernel, has a programming language of its own. The shell is a command interpreter. The shell’s work consists of programs and shell commands and utilities run by shell users. You can think of the shell running with the kernel as TSO/E running with MVS (after all, TSO/E is the original shell of MVS.)

A shell script consists of shell commands packaged in one file; a shell script is similar to an MVS command list or a REXX exec.

Accessing the Shell

There are several ways to access the shell. The most common are:

- The OMVS command, issued by a logged-on TSO/E user ID from a 3270 terminal or a workstation running a 3270 emulator, provides those who know TSO/E a familiar 3270-type interface.
- The telnet or rlogin commands, issued (without a login proc) from a workstation running TCP/IP.

When you use telnet or rlogin to access OS/390 UNIX, you have a standard UNIX interface, not a 3270-type interface. You work in character-at-a-time mode (each character processed as you type it). You can use the vi editor. One difference between telnet and login is:

- Communication through telnet can take place from OS/2, WIN95, VM, AIX, and many others. To login, type: telnet *hostname*
- Communication through rlogin takes place from another UNIX system. To login, type: rlogin *hostname*

The shell has some extended functions that have no MVS equivalents; for example:

- Pipes (directing output of a program to the input of another)
- Redirection (directing a data flow)

Daemons

A daemon is like an MVS started task — it is a long-running task that runs in the background and is not associated with any particular user. It starts work (or performs the work itself) on behalf of a user request, such as an rlogin request on an IP port, or a user shell script run at a scheduled time each day. Daemons usually run authorized and can issue authorized functions, such as changing the identity of the user associated with a process. You can think of the VTAM[®] started task as a daemon.

file; members of the owning group can read and execute the file, as can all users. The owner can write to the file; no one else can.

Superuser Authority

The installation defines certain system administrators as superusers, having UID(0), who can change the contents of any file, manage processes, and perform other administrative activities. When not doing activities that require superuser authority, system administrators can change to user authority, which permits access to his or her own files and other files, according to the permission bits.

Don't confuse superuser authority with MVS supervisor state. Being a superuser does not increase the ability to do things in MVS; being a superuser affects only your ability to do things in OS/390 UNIX. Remember the point about "no wall between MVS and UNIX"? Well, where security is concerned, the wall is there.

The Multiple Ways to Use OS/390 UNIX

IBM offers familiar ways to use OS/390 UNIX, whether your background has been with UNIX or with MVS. Workstation users can copy data from MVS data sets into the file system for use at the host or, through OS/390 Network File System, at the workstation. They can also upload ASCII-encoded files to the host, where they can be translated into EBCDIC encoding for the host. The same application program can access both traditional MVS data sets and files in the HFS. The program can invoke callable services from Assembler applications and subroutines, in addition to the C/C++ language applications. And, here are more ways:

- Interactive users can choose between the familiar interface of the UNIX shell or command interface and the traditional TSO/E panel or command interface. To use the HFS, you can use TSO/E commands, shell commands, or the panel interface of ISPF.

- To edit HFS files, you can use the ISPF full-screen editor or a UNIX editor, such as vi or ed, available in the shell.
- You can write MVS job control language (JCL) that includes UNIX shell commands.
- Using OS/390 UNIX extensions to REXX, you can run REXX programs from TSO/E, batch, the shell, or a C/C++ program.
- You can run a non-interactive shell command or a shell script from the TSO/E READY prompt and display the output at your terminal; you can run an interactive shell command through the OMVS command.
- You can use the BPXBATCH interface (with familiar JCL statements) to run the shell and its utilities in the background.

Workstation users can be connected to TSO/E and the shell through 3270 emulation or through OS/390 SecureWay Communication Server - IP (also known as TCP/IP). An installation can customize logging-on so that a user who prefers to work with a UNIX-like interface is put directly into the familiar UNIX shell. From there, the user can choose to switch to the TSO/E environment (through ISPF panels or TSO/E commands). An interactive user can log on through the TSO/E interface and switch to the shell when desired.

Comparing MVS and UNIX Concepts and Terms

Topics in the following table were collected from the MVS-OE mailing list, as mentioned earlier. The first column identifies a task or concept, in platform-independent terms; the second column describes the task or concept in MVS terms, the third column describes the task or concept in terms of AIX® (the UNIX that runs on RS/6000®), and the fourth column describes the task or concept for OS/390 UNIX. Note that the table is in random order; the disparate nature of the tasks and concepts in the table makes logical ordering and consistent phrasing difficult.

Comparing MVS, AIX and MVS Tasks and Concepts

Task or Concept	MVS	AIX	OS/390 UNIX
Part of the system that interacts with the processor	BCP (including the hardware configuration definition - HCD) and the DFSMSdfp™ element	The kernel	BCP (including the hardware configuration definition - HCD) and the DFSMSdfp element.
Virtual storage the system gives each user	MVS address space; each OS/390 user gets 2 gigabytes of virtual storage; some of that storage contains system code that is common for all users	Users get whatever virtual storage they need to reference, within the limits of the hardware and operating system	OS/390 UNIX work runs in MVS address spaces. In IBM manuals, "a UNIX address space" typically refers to an address space created (with WLM managing the initiators) by a fork or spawn function.
Data storage	Data sets	Files	Files
System configuration data	Parameters in members of SYS1.PARMLIB (also called parmlib) control how the system IPLs and how address spaces behave	The/etc file system controls characteristics. Also, the Object Data Manager (ODM) stores system configuration information	The/etc file system contains configuration files, product defaults and information that is used during system IPL and at other times, such as when daemons (tasks) are started. Unlike parmlib, it cannot be shared or concatenated, nor can it be shared among systems for searches purposes
Bit bucket	DD DUMMY	/dev/null file	/dev/null file
Locating data sets and files	Catalogs (master and user) and PDS directories	Directories	Directories
Locating file systems	The volume table of contents (VTOC) lists all data sets, but not members within the data sets	Because the file system is hierarchical, files can be located by the directories above them; the directories by the directories above them - right up to the root itself	Because the file system is hierarchical, files can be located by the directories above them; the directories by the directories above them - right up to the root itself
Shared libraries	Every address space can run a single copy of a program in the link pack area (LPA) for common use	Shared libraries are dynamic. When the first user loads a shared program, it is brought into storage. If no one is using the program, it can be freed	With OS/390 R9, shared libraries, similar to AIX. An HFS program that has the sticky bit on can invoke a program in the LPA

Comparing MVS, AIX and MVS Tasks and Concepts

Task or Concept	MVS	AIX	OS/390 UNIX
Text data representation	MVS's use of EBCDIC originated with the holes in punched cards. MVS does not care whether data is EBCDIC or ASCII. The applications themselves must handle the data, whatever the form. However, most MVS programs that work with data expect it in EBCDIC; an example is ISPF edit, which works only on EBCDIC data. When you use telnet to access TSO/E, the data you enter is converted from ASCII to EBCDIC before TSO/E processes it	AIX programs expect data to be in ASCII	OS/390 UNIX does not care whether data is EBCDIC or ASCII. All IBM-shipped OS/390 UNIX programs that work with data expect it in EBCDIC. Most programs that work with data in other UNIX systems produce it in ASCII. Therefore, if you port a program from another UNIX platform, you might have to convert the data, which you can do through the pax or iconv utilities. When you use telnet or rlogin to access the shell, the data you enter is converted from ASCII to EBCDIC before the shell processes it. When you use file transfer program (FTP) and do a text transfer on ASCII text, the file automatically converts to EBCDIC
Data Format	Record orientation, often an 80-byte record, reflecting the traditional punched card image	Byte orientation; organization of the data is provided by the application	Byte orientation; organization of the data is provided by the application
Data access methods	VSAM, BSAM, QSAM and others; also HFS	File organization is application's responsibility	File organization is application's responsibility; see also MVS
Case sensitivity	MVS can require uppercase (as in JCL and data set names); sometimes it folds to uppercase (as in commands and parmlib parameters)	Commands and file names are case sensitive	Commands and file names are case sensitive
Programming languages supported	COBOL, HLL Assembler, PL/I, C/C++, Fortran, APL2 [®] and others	C/C++, Java and MicroFocus COBOL	C/C++, Java, HLL Assembler, PL/I

Comparing MVS, AIX and MVS Tasks and Concepts

Task or Concept	MVS	AIX	OS/390 UNIX
Scripting languages	TSO/E CLISTS and REXX execs	Shell scripts, Perl, TCL or TK, awk	Shell scripts, REXX execs, awk and Perl
Online help	TSO/E HELP, ISPF help (PF1 key)	man (for manual) is a UNIX shell command that provides on-line help (called man pages) for shell commands and functions. Also, there is online help for Graphical User Interfaces and Task Guides	man provides on-line help for shell commands and functions. Also, OHELP (an interface to Book Manager) is a UNIX shell command that provides on-line help for the TSO/E UNIX commands
Smallest element that performs work	A task (represented by an MVS control block called a TCB). MVS supports multiple tasks	A thread. The kernel supports multiple threads. When users start a unit of work, they are executing a process, and functions like exec and fork create new processes. A process can have several threads	A thread is implemented as a task. The kernel supports multiple threads in an address space. When users start a unit of work, they are executing a process, and functions like exec and fork create new processes. A process can have several threads
A long-running unit of work	A started task; sometimes a long-running job	A daemon	A daemon or a started task
Order in which the system searches for programs to run	The search order for programs in libraries is usually: TASKLIB, STEPLIB, JOBLIB, LPALST and LNKLST	Programs are loaded from the file system according to the user's PATH environmental variable (a list of directories to be searched). For dynamic link library (DLL) programs, programs are loaded from LIBPATH	Programs are loaded from the file system according to the user's PATH environmental variable. If the kernel encounters a program with the sticky bit enabled, the search order reverts to the MVS search order

Comparing MVS, AIX and MVS Tasks and Concepts

Supplement - UNIX on OS/390

Task or Concept	MVS	AIX	OS/390 UNIX
Backups	Backups are done concurrently in the background on a single instance of the data, while other work continues to run with the required response times and service levels. Data being archived cannot change during the archive process	An AIX system can mirror at the logical disk volume level. If a disk fails, it can be removed and a new disk added. This method is only available locally. Another backup technique is tools that copy data to another media, usually tape. The tools require that data being archived not change during the archive process.	Backups are done concurrently in the background, on a single instance of the data, while other work continues to run with the required response times and service levels. Data being archived cannot change during the archive process
Backup and recovery utilities	Hierarchical Storage Manager (DFSMSHsm™) and DFSMSdss™ elements; IBM product Automated Data Storage Management (ADSM)	The pax, tar and cpio utilities	The pax, tar, and cpio utilities and ADSM
Availability of the system	Unscheduled outages are rare. Systems generally run fully available. When backups, database reorganizations, and other software housekeeping occur, they typically run concurrently alongside diverse production workloads, such as online processing, query and webserving.	Unscheduled outages tend to be more frequent on UNIX systems; the outages are tolerable if a reboot takes a short time and if backup and recovery procedures are not required to ensure the integrity of data. Reorganizations of databases might require a scheduled outage.	Unscheduled outages are rare. Systems generally run fully available. When backups, database reorganizations, and other software housekeeping occur, they typically run concurrently alongside diverse production workloads, such as online processing, query and webserving.
Assigning physical storage	Allocate storage using dynamic allocation (through ISPF), the TSO/E ALLOCATE command, and the DD statement in JCL	A system administrator assigns file systems to logical volumes that are mapped to physical storage	A system administrator allocates HFS data sets associated with the file systems

Comparing MVS, AIX and MVS Tasks and Concepts

Task or Concept	MVS	AIX	OS/390 UNIX
Accounting for use of system resources	SMF (system measurement facility) or a non-IBM product, at an address space level	AIX provides system accounting utilities	SMF
Balancing workload	Workload manager (WLM) policy or settings in parmlib members IEAICSxx and IEAIPSxx	AIX uses process priorities that can be assigned by the nice command	A WLM policy or settings in IEAIPSxx
Monitoring performance	The RMF element	Performance tools check the hardware resources, CPU utilization, memory, I/O throughput and disks	Non-IBM vendor products provide monitoring tools that allow you to see what is going on and take actions against user processes. RMF only provides operator commands that allow you to view output on the system console
Determining problems	TSO TEST or the IBM Debug Tool that comes with the C/C++ element. MVS has console messages, a hardcopy log, dumps and abend codes. You can set SLIP traps; also, you can use generalized trace facility (GTF) and component trace. With JES2, use SDSF to view the system log	Ertnos, visual source-level debuggers, syslog, core dumps; also dbx symbolic debugger and errpt for system errors, and core dumps	Ertnos and ertrjrs, which are return codes and reason codes, the dbx debugger, or VA Debugger for C/C++ and Java. The syslog daemon SYSLOGD creates the UNIX log, which you can browse. Set SLIP traps to get dumps, identifying which jobname to set the trap for; component trace and LE dumps and utilities
Managing product installation and service	The SMP/E element manages product installation and service	AIX system commands manage administrative operations. The system management interface tool (SMIT) is a menu interface to some of these commands	SMP/E manages product installation and service

Comparing MVS, AIX and MVS Tasks and Concepts

Task or Concept	MVS	AIX	OS/390 UNIX
Using the system interactively	Users log on to the system through TSO/E; the ISPF element provides a panel-driven interface to TSO/E. A userid can have only one logon session active at a time	Users login to systems and execute shell sessions in the shell environment. They can issue the rlogin or telnet commands to connect to an operating system. Additionally, AIX users have a graphical interface called aixterm. AIX users can have many logins open at once	OS/390 userids can use the rlogin and telnet commands and have multiple login sessions. Additionally, they can log on to TSO/E and enter the OMVS command
Starting a program to run in the background	Submit a batch job to run through the job's JCL	Shell commands that use the "&" parameter; cron command	Shell commands that use the "&" parameter; cron, and the BPXBATCH interface to submit shell commands and utilities
Scheduling a program	JES2, JES3, System Automation for OS/390	cron, programs (typically daemons) started at boot time, which are listed in /etc/inittab; also, the System Resource Controller (SRC) starts and restarts daemons	cron or /etc/rc or another facility for starting and restarting daemons
Starting a program to run in the foreground	You can use ISPF or TSO/E to run a job in the foreground	Issue the name of the shell script or program	Issue the name of the shell script or program
Starting a process	Use an ATTACH or XCTL macro, invoke a CALL from an already-running program, EXECUTE in TSO/E, or submit a batch job	fork/exec functions and p_thread create	fork/exec functions and p_thread create
Creating a new address space	ASCRE macro, logging on to TSO/E, issuing the START command, submitting a batch job, APPC, the scheduler, and WLM	AIX does not have the same concept of address spaces as MVS does	fork, sometimes spawn functions, or a login command start an address space
Creating an executable program	Program binder and compilers	makefiles, program binders, compilers, gnu, emacs, lex, yacc, make	makefiles, program binders, compilers, lex, yacc

Comparing MVS, AIX and MVS Tasks and Concepts

Task or Concept	MVS	AIX	OS/390 UNIX
Editing data or code	ISPF editor	Many editors exist, such as vi, ed, sed and emacs	vi, ed, oedit (an ISPD editor) and sed: <ul style="list-style-type: none"> • Those coming from TSO/E would probably use oedit (in OMVS under TSO/E and ISPF). • Those coming from UNIX could use vi (only from rlogin or telnet sessions). To convert ASCII/EBCDIC, download and use nedit, an X Windows-based editor with a GUI interface. Emacs and Visula slickedit are available from vendors
Managing programs	For JES2 users, the SDSF element allows you to view jobs and purge jobs	The ps shell command allows you to view processes and threads, and kill jobs. The kill command terminates a process. To terminate an interactive process, use ctrl-C to send a SIGKILL interrupt to the process	The ps shell command allows you to view processes and threads, and kill jobs. The kill command terminates a process. To terminate an interactive process, use ctrl-C to send a SIGKILL interrupt to the process. Also, SDSF
Authorizing access to do certain functions	FACILITY classes, in RACF	Through the execute permission bit with UID and GID definitions. Also, access control lists through security mechanisms, such as Kerberos and DCE	Through permission bits, UIDs and GIDs. Use of UNIXPRIV FACILITY class grants certain privileges. There are no access control lists
Serializing use of data	Global resource serialization, which processes RESERVE or ENQ/DEQ macros; latches and locks; compare and swap operations	byte-range locking of files	byte-range locking of files
Changing from unauthorized to authorized	supervisor call (SVC), PC instruction, or MODESET macro	setuid bit	setuid bit

Notes & Doodles

Supplement - UNIX on OS/390

This supplement is the beginning of a larger work on UNIX and OS/390, which will be published on the Web in the future. If you are interested in providing input, please send a note to duna@us.ibm.com

Notes & Doodles

This supplement is the beginning of a larger work on UNIX and OS/390, which will be published on the Web in the future. If you are interested in providing input, please send a note to duna@us.ibm.com



HOT TOPICS

**An OS/390 Newsletter for Release 9
February 2000**

Executive Editor Bob Ward

Managing Editor Dick Wagenaar

Art Director Gene Posca

Graphic Designer Stacy Newman

© International Business Machines
Corporation, 1999, 2000

Produced in the United States of America
2-2000
All Rights Reserved

The OS/390 Hot Topics Newsletter is published twice yearly as a service to OS/390 customers, providing articles of interest to the OS/390 community.

The information in this Newsletter is presented as is and is based on the best knowledge of the authors. No warranty is provided (neither expressed nor implied).

IBM may not offer the products, features or services discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM representative for information on the products or services available in your area.

© IBM, AIX, APL2, BookManager, DB2, ESCON, Language Environment, OpenEdition, OS/2, OS/390, Parallel Sysplex, RACF, RS/6000, S/390, SecureWay, VTAM are registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

™ Distributed Relational Database Architecture, DFSMSdftp, DFSMSHsm, DFSMSdss, MVS, RMF are trademarks of International Business Machines Corporation in the United States, other countries, or both.

Java, and all Java-based trademarks are trademarks of Sun Microsystems, Inc., in the United States, or other countries, or both.

Microsoft, Windows, Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark in the United States and other countries licensed exclusively through The Open Group.

Other company, product, and service names may be trademarks or service marks of others.



GA22-7431-01