26 RACF

November 2001

update

In this issue

RACF Update

Published by

Xephon
27-35 London Road
Newbury
Berkshire RG14 1JL
England
Telephone: 01635 38030
From USA: 01144 1635 38030
E-mail: fionah@xephon.com

North American office

Xephon
Post Office Box 350100
Westminster CO 80035-0100
USA
Telephone: (303) 410-9344

RACF Update on-line

Code from *RACF Update*, and complete issues in Acrobat PDF format, can be downloaded from http://www.xephon.com/racf; you will need to supply a word from the printed issue.

Subscriptions and back-issues

A year's subscription to *RACF Update* (four quarterly issues) costs £190.00 in the UK; $290.00 in the USA and Canada; £196.00 in Europe; £202.00 in Australasia and Japan; and £200.50 elsewhere. The price includes postage. Individual issues, starting with the August 1995 issue, are available separately to subscribers for £48.50 ($72.75) each including postage.

Editor

Fiona Hewitt

Disclaimer

Readers are cautioned that, although the information in this journal is presented in good faith, neither Xephon nor the organizations or individuals that supplied information in this journal give any warranty or make any representations as to the accuracy of the material it contains. Neither Xephon nor the contributing organizations or individuals accept any liability of any kind howsoever arising out of the use of such material. Readers should satisfy themselves as to the correctness and relevance to their circumstances of all advice, information, code, JCL, and other contents of this journal before making any use of it.

Contributions

When Xephon is given copyright, articles published in *RACF Update* are paid for at £170 ($260) per 1000 words and £100 ($160) per 100 lines of code for the first 200 lines of original material. The remaining code is paid for at the rate of £50 ($80) per 100 lines. In addition, there is a flat fee of £30 ($50) per article. To find out more about contributing an article, without any obligation, please contact us at any of the addresses above or download a copy of our *Notes for Contributors* from www.xephon.com/nfc

Printed in England.

Confirming a user's RACF access to generation datasets residing on tape

The IFG0EX0B exit presented here is invoked during OPEN's initial processing of all datasets that are located on tape and direct access devices. This version of IFG0EX0B processes only those OPEN requests that are directed toward generation datasets that reside on tape volumes. It's intended for use by installations that are using the OS/Security Server and which do not have TAPEDSN activated.

Whenever an attempt to OPEN such a dataset is made by a program, a RACROUTE macro is issued, to ascertain whether or not the user who submitted the job can legitimately access the dataset being OPENed. The ACEE operand is not used on the RACROUTE macro, so RACF will use either the task's ACEE or the address space's ACEE to check the user's authorization to access a particular dataset. If he or she has access to the dataset, processing of the OPEN request is allowed to continue; if not, an error message that contains his/her name and the name of the dataset that (s)he's attempting to access is displayed on consoles in the computer operations area.

Below is an example of the messages that would have been issued if a user named Fiona Smith had not previously been granted access to a dataset named XEPHON.RACF.G0001V00:

```
ICH408I USER(USER1 ) GROUP(BA000   ) NAME(FIONA. SMITH       )
  XEPHON.RACF.G0001V00 CL(DATASET ) VOL(111111 )
  INSUFFICIENT ACCESS AUTHORITY
  FROM XEPHON.RACF.* (G)
  ACCESS INTENT(READ   )  ACCESS ALLOWED(NONE   )
FIONA SMITH
HAS NO RACF AUTHORIZATION TO ACCESS
XEPHON.RACF.G0001V00
```

Afterwards, the task that issued the OPEN request is abnormally terminated with an abend code of S913, along with a reason code that matches the return code from RACROUTE.

If Fiona should have had access to the dataset, a RACF PERMIT DATA SET with READ authority must first be issued to grant her the necessary access.

A SETR REFRESH GENERIC(DATASET) command would also have to be issued since this is a generic profile for a dataset.

A word of warning: all UCBs are assumed to be located below the 16M line. Also, it's assumed that there will always be 1,000 bytes of free storage available for IFG0EX0B's use – that is, an unconditional STORAGE request of that amount may be requested during every OPEN request.

Only a return code of eight is considered relevant by this version of IFG0EX0B, since protecting datasets via RACF is not a requirement in my shop. For installations that do require RACF protection of all datasets, IFG0EX0B must be modified to consider return codes other than zero and eight.

SMP/E CONTROL STATEMENTS

The following SMP/E control statements were used to install IFG0EX0B into SYS1.LPALIB. An IPL with a CLPA was required to implement it. The reason that a REJECT control statement is present is that my code failed the first time that I used it, so another attempt was required before successful results were achieved. Imagine that!

Only a few of the DD statements that were used in SMP processing are provided, since so many of them are site-dependent. Note that IFG0EX0B must first have been assembled with its generated object code stowed in a dataset named NTUSER.LROBJ for the following set-up to work properly.

The FMID used in this example is for OS/390 release 2.9. For OS/390 release 2.6, use HDZ11D0. (I don't know what to use for other releases.)

```
//IFGØEXØB EXEC PGM=GIMSMP,REGION=6M
//LPALIB DD DISP=SHR,DSN=SYS1.LPALIB
...
//USEROBJ  DD   DISP=SHR,DSN=NTUSER.LROBJ(IFGØEXØB)
//SMPCNTL  DD   *
 SET BOUNDARY(GLOBAL) .
 REJECT  S(NTØØØØ1)  BYPASS(APPLYCHECK).
 RECEIVE S(NTØØØØ1) SOURCEID(LOCALMOD) .
 SET BOUNDARY(MVST1ØØ) .
 APPLY S(NTØØØØ1) REDO.
```

```
//SMP.SMPPTFIN DD  *
++USERMOD(NT00001) .
++VER(Z038) FMID(HDZ11E0) .
++MOD(IFG0EX0B) TXLIB(USEROBJ).
```

IFG0EX0B

```
TITLE 'INSTALLATION-WIDE OPEN EXIT - IFG0EX0B'
*********************************************************************
*                                                                   *
*        THIS VERSION OF IFG0EX0B CONFIRMS A USER'S RACF ACCESS     *
*        TO GENERATION DATA SETS THAT RESIDE ON TAPES.              *
*        IFG0EX0B IS EMBEDDED WITHIN IGC0001I IN SYS1.LPALIB.       *
*        THE ACEE THAT IS ASSOCIATED WITH THE TASK OR ADDRESS SPACE *
*        INVOLVED IN AN OPEN REQUEST IS USED FOR RACROUTE REQUESTS. *
*                                                                   *
*        CONTENTS OF REGISTERS AT ENTRY TO IFG0EX0B ARE AS FOLLOWS: *
*        R1  - ADDRESS OF IFG0EX0B'S PARAMETER LIST                 *
*        R13 - ADDRESS OF AN 18-WORD SAVE AREA                      *
*        R14 - RETURN ADDRESS - NATURALLY                           *
*        R15 - IFG0EX0B'S ENTRY POINT                               *
*        CONTENTS OF ALL OTHER REGISTERS - IRRELEVANT               *
*                                                                   *
*        UNLESS IT ISSUES AN S913 ABEND, THIS EXIT ALWAYS DEPARTS   *
*        WITH A RETURN CODE OF ZERO.                                *
*                                                                   *
*********************************************************************
         SPACE 3
         MACRO
&TAPNAME TAPINFO
         DS    0F
         PUSH PRINT
         PRINT GEN
&TAPNAME DC    CL8'&SYSECT'
         DC    A(&SYSECT)
         DC    CL6'&SYSTIME'
         DC    CL8'&SYSDATE'
         POP  PRINT
         MEND
         EJECT
IFG0EX0B CSECT
IFG0EX0B AMODE 24
IFG0EX0B RMODE 24
         PRINT NOGEN
         SPACE
         USING IFG0EX0B,R12        ESTABLISH IFG0EX0B ADDRESSABILITY
         USING OIEXL,R1            ESTABLISH OIEXL ADDRESSABILITY
         USING PSA,R0              ESTABLISH PSA ADDRESSABILITY
         SPACE
```

```
         BAKR  R14,RØ              PRESERVE ENVIRONMENT AT ENTRY-TIME
         LR    R12,R15             PRIME BASE REGISTER
         SR    R13,R13             INDICATE NO WORKING STORAGE PRESENT
         SPACE
         ICM   R1Ø,15,OIEXTIOT     POINT TO TIOT ENTRY
         BE    PPGDUST             EXIT IF UNAVAILABLE-SHOULD NEVER BE
         USING TIOENTRY,R1Ø        ESTABLISH TIOENTRY ADDRESSABILITY
         DROP  R1                  FORGET OIEXL
         SPACE
         SR    R9,R9               ZERO INDEX REGISTER
         LR    R2,R9               CLEAR UCB REGISTER
         LR    R8,R9               CLEAR JFCB REGISTER
         SPACE
PPGHAVIT ICM   R2,7,TIOEFSRT       POINT TO UCB
         BE    PPGNXTDD            BRANCH IF POINTER IS ABSENT
         USING UCBOB,R2            ESTABLISH UCB ADDRESSBILITY
         TM    UCBTBYT3,UCB3TAPE   TEST IF THIS IS A TAPE DEVICE
         BNO   PPGNXTDD            IF NOT, PROCESS NEXT TIOT ENTRY
         SPACE
         ICM   R8,7,TIOEJFCB       POINT TO JFCB PREFIX
         BE    PPGNXTDD            BRANCH IF POINTER IS ABSENT
         SPACE
         LA    R8,16(R8)           POINT TO JFCB PROPER
         USING INFMJFCB,R8         ESTABLISH JFCB ADDRESSABILITY
         TM    JFCBIND1,JFCGDG     TEST FOR A GENERATION DATA SET
         BO    PPGDGDSN            BRANCH IF 'TIS GDG
         SPACE
         DROP  R2                  FORGET UCB
         EJECT
*********************************************************************
*                                                                   *
*        DETERMINE IF A DATA SET IS A GDG THE HARD WAY.  BEGIN AT ITS *
*        END AND LOCATE THE FIRST DELIMITER.  A DATA SET IS ASSUMED   *
*        TO BE A GDG IF ITS LOW-LEVEL QUALIFIER IS OF THE FORM        *
*        .G....V.. .                                                  *
*                                                                   *
*********************************************************************
         SPACE
         LA    R1,JFCBDSNM+43      POINT TO END OF NAME OF DATA SET
         LA    RØ,42               SET MAXIMUM LENGTH OF NAME
PPGETNAM CLI   Ø(R1),C' '          TEST IF END OF NAME
         BNE   PPGOTNAM            BRANCH IF SO
         BCTR  R1,RØ               POINT TO PREVIOUS CHARACTER
         BCT   RØ,PPGETNAM         LOOP POWER!
         B     PPGNXTDD            DO NEXT DD IF NO NAME
         SPACE
PPGOTNAM CLI   Ø(R1),C'.'          TEST FOR A PERIOD
         BE    PPGCKNAM            BRANCH IF SO
         BCTR  R1,RØ               POINT TO PREVIOUS CHARACTER
         BCT   RØ,PPGOTNAM         CONTINUE SEARCHING FOR AN EPOCH
```

```
         B     PPGNXTDD            DO NEXT DD IF NOT GDG
         SPACE
PPGCKNAM CLI   1(R1),C'G'          TEST IF BEGINNING FORMAT OF A GDG
         BNE   PPGNXTDD            BRANCH IF NOT
         CLI   6(R1),C'V'          TEST IF ENDING FORMAT OF A GDG
         BNE   PPGNXTDD            BRANCH IF NOT
         CLI   7(R1),C' '          TEST IF ENDING FORMAT OF A GDG
         BE    PPGNXTDD            BRANCH IF SO
         CLI   8(R1),C' '          TEST IF ENDING FORMAT OF A GDG
         BE    PPGNXTDD            BRANCH IF SO, ELSE ASS-U-ME GDG
         EJECT
*********************************************************************
*                                                                   *
*        ACQUIRE WORKING STORAGE THEN, WITHIN IT,                   *
*        CONSTRUCT A PARAMETER LIST FOR RACROUTE.                   *
*                                                                   *
*********************************************************************
         SPACE
PPGDGDSN LTR   R13,R13             TEST FOR PRESENCE OF WORKING STORAGE
         BNE   PPGGSTOR            BRANCH IF PRESENT
         SPACE
         LA    R5,PPGLEND          SET LENGTH OF WORK AREA REQUIRED
         STORAGE OBTAIN,LENGTH=(5) ACQUIRE VIRTUAL STORAGE
         LR    R13,R1              POINT TO SAVE AND WORK AREAS
         USING PPGWORK,R13         ESTABLISH PPGWORK ADDRESSABILITY
         SPACE
PPGGSTOR LA    RØ,44               SIZE OF THE NAME OF A DATA SET
         STH   RØ,PPGBUFLN         SET LENGTH OF BUFFER CONTAINING NAME
         STCM  RØ,12,PPGDSNLN      SET MAX LENGTH OF NAME TO ZERO
         MVC   PPGDSN,JFCBDSNM     SET NAME OF DATA SET IN PARM AREA
         LA    R3,JFCBVOLS         POINT TO FIRST VOLUME SERIAL NUMBER
         DROP  R8                  FORGET JFCB
         SPACE
         LA    RØ,PPGCOMM          POINT TO WORK AREA FOR RACF
         LA    R1,PPGRACM          POINT TO 'L' FORM OF RACROUTE MACRO
         MVC   PPGRACM(PPGRACRL),PPGRACR PRIME MACRO'S DYNAMIC AREA
         SPACE
         RACROUTE REQUEST=AUTH,                                    C
               ENTITYX=(PPGBUFLN),                                 L
               VOLSER=(R3),                                        A
               RELEASE=1.9,                                        I
               WORKA=(Ø),                                          R
               ATTR=READ,                                          E
               DSTYPE=N,                                           M
               MF=(E,(1))
         SPACE 1
         LTR   R7,R15              TEST RETURN CODE
         BE    PPGNXTDD            BRANCH IF ACCESS TO DSN BY USEROK
         C     R7,PPGF8            TEST IF RETURN CODE OF EIGHT
         BE    PPGNOGO             BRANCH IF NO ACCESS TO DSN BY USER
```

```
         SPACE 1
PPGNXTDD IC    R9,TIOELNGH         LENGTH OF DD ENTRY
         LA    R1Ø,Ø(R9,R1Ø)         NEXT DD ENTRY
         SPACE 1
PPGTEND  SR    R1,R1               CLEAR REGISTER FOR COMPARE
         C     R1,TIOENTRY         TEST IF END OF TIOT
         BE    PPGCIAO             BRANCH IF SO
         CLC   TIOEDDNM,PPGCLN     TEST FOR DDNAME OF BLANKS
         BNE   PPGCIAO             EXIT IF NOT
         B     PPGHAVIT            PROCESS CONCATENATED DD STATEMENT
         EJECT
*********************************************************************
*                                                                   *
*        CONSTRUCT THEN ISSUE MESSAGE INDICATING THAT ACCESS TO     *
*        A DATA SET BY A USER WAS DENIED.  AFTERWARDS ABRUPTLY      *
*        TERMINATE THE CURRENT TASK WITH AN ABEND CODE OF S913.     *
*                                                                   *
*********************************************************************
         SPACE 1
PPGNOGO  LA    RØ,L'PPGNOCON       MAXIMUM LENGTH OF INFORMATIVE CON-
         STH   RØ,PPGWLS2L          STANT TO WTO AREA
         MVC   PPGWLS2,PPGNOCON    MOVE CONSTANT TO WTO AREA
         SPACE 1
         LA    RØ,44               MAXIMUM LENGTH OF NAME OF DATA SET
         STH   RØ,PPGWLS3L          TO WTO AREA
         MVC   PPGWLS3,PPGDSN      STOW NAME OF DATA SET IN WTO AREA
         SPACE 1
         L     R8,PSAAOLD          RETRIEVE ADDRESS OF ASCB
         USING ASCB,R8             ESTABLISH ASCB ADDRESSABILITY
         SPACE 1
         LA    RØ,3Ø               LENGTH OF NAME OF USER - MAXIMUM
         STH   RØ,PPGWLS1L          TO WTO AREA
         MVC   PPGWLS1,=CL3Ø'UNKNOWN USER'
         SPACE 1
         L     R1,ASCBASXB         ADDRESS SPACE BLOCK EXTENSION
         USING ASXB,R1
         ICM   R2,15,ASXBSENV      ADDR ACCESS CNTL ENVIRONMENT ELEMENT
         BE    PPHDOWTO            BRANCH IF NONEXISTENT
         DROP  R1                  FORGET ASXB
         SPACE
         SR    R8,R8               CLEAR WORK REGISTER
         USING ACEE,R2
         ICM   R3,15,ACEEUNAM      FETCH ADDRESS OF NAME OF USER
         BZ    PPHDOWTO            IF NONEXISTENT, ISSUE WTO
         SPACE
         ICM   R8,1,Ø(R3)          RETRIEVE LENGTH OF USER'S NAME
         BZ    PPHDOWTO            BRANCH IF NONEXISTENT
         SPACE
         CH    R8,PPHH3Ø           TEST IF WITHIN LIMITS
         BL    PPHLOK              BRANCH IF SO
```

```
         LH    R8,PPHH3Ø           SET UPPER BOUNDS
PPHLOK   STH   R8,PPGWLS1L         REVISE LENGTH OF USER'S NAME
         BCTR  R8,RØ               REDUCE BY ONE FOR EX INSTRUCTION
         EX    R8,PPHMNAME         STOW NAME IN WTO AREA
         SPACE 1
PPHDOWTO LA    R3,PPGWLS1L         POINT TO NAME OF USER INFORMATION
         LA    R4,PPGWLS2L         POINT TO CONSTANT INFORMATION
         LA    R5,PPGWLS3L         POINT TO DATA SET INFORMATION
         SPACE 1
         MVC   PPGWTONB(PPGWTLNB),PPGWTOLS WTO PATTERN TO OUTPUT AREA
         LA    R1,PPGWTONB         POINT TO LIST FORM OF WTO MACRO
         WTO   TEXT=(((3),),((4),),((5),)),MF=(E,(1)) INFORMATIVE MSG
         EJECT
*********************************************************************
*                                                                   *
*        TERMINATE THIS ACTIVITY - HERE AND NOW!                    *
*                                                                   *
*********************************************************************
         SPACE 1
         L     1,PPG913            SET RACF ABEND CODE OF 913
         ICM   R1,8,PPGXØ4         REASON CODE SPECIFIED; NO DUMPOPT
         LR    15,7                SET REASON CODE
         SVC   13                  TERMINATE THIS TASK WITH S913
         SPACE 3
*********************************************************************
*                                                                   *
*        CLEAN-UP AND THEN TERMINATE                                *
*                                                                   *
*********************************************************************
         SPACE 1
PPGCIAO  LTR   R13,R13             TEST IF WORKING STORAGE ACQUIRED
         BE    PPGDUST             BRANCH IF NOT
         SPACE 1
         LA    R5,PPGLEND          SET LENGTH OF WORK AREA OBTAINED
         STORAGE RELEASE,LENGTH=(5),ADDR=(R13) FREE IT
         SPACE
PPGDUST  SR    R15,R15             SET ZERO FOR A RETURN CODE
         PR    R14                 BACK TO DUST
         EJECT
*********************************************************************
*                                                                   *
*        CONSTANTS AND OTHER SUCH NONSENSE                          *
*                                                                   *
*********************************************************************
         SPACE
         TAPINFO
         SPACE
PPHMNAME MVC   PPGWLS1(*-*),1(R3)  ====>  EXECUTE ONLY  <====
PPHH3Ø   DC    H'3Ø'
PPGXØ4   DC    X'Ø4'
```

```
PPGCLN   DC    CL8' '
         SPACE
PPGWTOLS WTO   TEXT=((,C),(,D),(,DE)),MF=L
PPGWTLNB EQU   *-PPGWTOLS
         SPACE
PPGNOCON DC    C'HAS NO RACF AUTHORIZATION TO ACCESS'
         SPACE
PPGF8    DC    F'8'
PPG913   DC    XL4'00913000'
         SPACE
PPGRACR  RACROUTE REQUEST=AUTH,                                        M
               CLASS='DATASET',                                        E
               MF=L
PPGRACRL EQU   *-PPGRACR
         SPACE
PPG80    DC    X'80'
         EJECT
*************************************************************************
*                                                                       *
*        DSECT FOR WORK-AREA USED TO ISSUE RACROUTE AND WTO MACROS      *
*                                                                       *
*************************************************************************
         SPACE
PPGWORK  DSECT
         DS    18F
         SPACE
PPGRACM  DS    XL(PPGRACRL)
         DS    0F
PPGBUFLN DS    XL2
PPGDSNLN DS    XL2
PPGDSN   DS    CL44
         DS    0D
PPGCOMM  DS    CL512
         SPACE
         DS    0F
PPGWTONB DS    CL(PPGWTLNB)
         SPACE
         DS    0F
PPGWLS1L DS    XL2
PPGWLS1  DS    CL30
         DS    0F
PPGWLS2L DS    XL2
PPGWLS2  DS    CL(L'PPGNOCON)
         DS    0F
PPGWLS3L DS    XL2
PPGWLS3  DS    CL44
         SPACE
PPGLEND  EQU   *-PPGWORK
         TITLE 'IFG0EX0B - GENERATE OS/390 CONTROL BLOCKS'
         YREGS
```

```
         SPACE
         IHAACEE
         SPACE
         IHAPSA
         SPACE
         IHAASCB
         SPACE
         IHAASXB
         SPACE
         IECOIEXL DSECT=YES
         SPACE
         DSECT
         IEFJFCBN
         SPACE
PPGUCB   DSECT
         IEFUCBOB ,
         SPACE
PPGTIOT  DSECT ,
         IEFTIOT1
         SPACE 1
         END
```

Julia H Pond
Information Systems Technical Consultant (USA)

RACF operations attribute for storage managers – for and against

Storage managers have to deal with data which is RACF protected and for which they are not explicitly authorized (in RACF terms). In the past, there were two ways to enable storage managers to do their work:

- The most popular method is to use the OPERATIONS attribute. This generally grants ALTER access to all datasets except those for which an OPERATIONS user is not explicitly excluded.

  However, this attribute may grant too much authority. Accesses granted by the OPERATIONS attribute can be audited using the SETROPTS(OPERAUDIT) option. But remember that SMF records do not show whether the dataset has just been moved to another volume, has been dumped to tape, or has really been read or updated by the OPERATIONS user.

- A second way is to use ALTER access in class DASDVOL. However, this cannot be audited and is not generally sufficient.

Some years ago, DFSMSdss introduced a third method, whereby a series of FACILITY class profiles can be used to control access to specific types of DFSMSdss processing (PGM=ADRDSSU).

This article considers the ways in which you can enable storage managers to do their work, while at the same time removing as many OPERATIONS attributes as possible, and ensuring that everyone has only the lowest levels of access necessary for them to perform their jobs.

COMMON TYPES OF WORK

First, let's have a look at the common types of work of storage administrators, and consider the required access authorizations. There are two major objects to deal with, namely whole DASD volumes and single datasets.

- DASD volumes can be:

  - initialized (1)
  - dumped (2) (Note that in this case ‘dump’ is taken to mean the same as back-up.)
  - restored (3)
- Datasets can be:
  - dumped (4)
  - restored (5)
  - copied (6)
  - deleted (7)
  - migrated (8)
  - recalled (9)
  - read (10)
  - updated (11).

In the sections below, I list the tools required for these tasks, and identify the required authorizations. Note, however, that because tasks (10) and (11) are outside the scope of work of a storage administrator, they’re not discussed in this article.

INITIALIZING A DASD VOLUME (TASK 1)

Tool

- ICKDSF – batch job.

Required authorizations

- Volume has to be OFFLINE[1]
- Operator-Reply[1]
- ALTER authority in DASDVOL class.

([1] These are not authorizations in RACF terms, but they are used to control this type of access.)

DUMP A DASD VOLUME (TASK 2)

Tool

- ADRDSSU – batch job.

Required authorizations

- ALTER authority in DASDVOL class.

or

- READ authority for STGADMIN.ADR.STGADMIN.DUMP profile in FACILITY class, and
- The use of the ADRDSSU keyword ADMINISTRATOR.

Sample SYSIN

```
//SYSIN     DD *
  DUMP INDD(DASD) -
       OUTDD(CART) -
       OPTIMIZE(4) -
       ALLEXCP     -
       ADMINISTRATOR    -
       ALLDATA(*)  -
       COMPRESS
/*
```

RESTORE A DASD VOLUME (TASK 3)

Tool

- ADRDSSU – batch job.

Required authorizations

- ALTER authority in DASDVOL class.

or

- READ authority for STGADMIN.ADR.STGADMIN.RESTORE profile in FACILITY class, and the use of the ADRDSSU keyword ADMINISTRATOR.

Sample SYSIN

```
//SYSIN     DD *
RESTORE FULL INDD(CART) -
OUTDD(DASD) -
          ADMINISTRATOR -
          COPYVOLID -
PURGE -
 WAIT(2,2)

/*
```

DUMP A DATASET (TASK4 , METHOD 1)

Tool

- ADRDSSU – batch job.

Required authorizations

- ALTER authority in DASDVOL class (NONSMS only).

or

- READ authority for STGADMIN.ADR.STGADMIN.DUMP profile in FACILITY class, and the use of the ADRDSSU keyword ADMINISTRATOR.

Sample SYSIN

```
//SYSIN     DD *
DUMP DATASET(INCLUDE(HLQ.SAMPLE*.**)) -
     OUTDDNAME(CART) -
     OPTIMIZE(4) -
     ALLEXCP     -
     ALLDATA(*)  -
     ADMINISTRATOR -

     COMPRESS

/*
```

RESTORE A DATASET (TASK 5, METHOD 1)

Tool

- ADRDSSU – batch job.

Required authorizations

- ALTER authority in DASDVOL class (NONSMS only).

or

- READ authority for STGADMIN.ADR.STGADMIN.RESTORE profile in FACILITY class, and the use of the ADRDSSU keyword ADMINISTRATOR.

Sample SYSIN

```
//SYSIN     DD *
DUMP DATASET(INCLUDE(HLQ.SAMPLE*.**)) -
     OUTDDNAME(CART) -
     OPTIMIZE(4) -
     ALLEXCP     -
     ALLDATA(*)  -
     ADMINISTRATOR -

     COMPRESS

/*
```

DUMP A DATASET (TASK 4, METHOD 2)

Tool

- DFSMShsm (HBACKDS).

Required authorizations

- AUTH userid DATABASEAUTHORITY(USER).

Sample command

- HBACKDS dataset.name

RESTORE A DATASET (TASK 5, METHOD 2)

Tool

- DFSMShsm (HRECOVER).

Required authorizations

- AUTH userid DATABASEAUTHORITY(USER).

Sample command

- HRECOVER dataset.name [REPLACE]

COPY A DATASET WITH RENAME (TASK 6)*

Tool

- ADRDSSU – batch job

Required authorizations

- READ authority for STGADMIN.ADR.STGADMIN.COPY.RENAME profile in FACILITY class.
- The use of the ADRDSSU keyword ADMINISTRATOR.

Sample SYSIN

```
//SYSIN     DD *
COPY DATASET(INCLUDE(hlq.sample.**)) -
     OUTDDNAME(DASDOUT) -
     ADMINISTRATOR -
RENUNC((hlq.sample.**,hlq.example.**))

/*
```

DFSMS does not allow duplicate uncatalogued datasets. By adding a delete step, this scenario can also be used to perform a ‘rename’.

DELETE A DATASET (TASK 7)

Tools

- IEFBR14 – batch job
- ISPF 3.*x*
- Any other delete as IDCAMS, TSO...

Required authorizations

- ALTER authority in DASDVOL class (NONSMS only) (dataset will be deleted, but not uncatalogued).

or

- ALTER authority to the catalog.

MIGRATE A DATASET (TASK 8)

Tool

- DFSMShsm.

Required authorizations

- AUTH userid DATABASEAUTHORITY(USER).

RECALL A DATASET (TASK 9)

Tool

- DFSMShsm.

Required authorizations

- AUTH userid DATABASEAUTHORITY(USER).

SUMMARY

In order to perform the tasks discussed above, you need one of following tools:

- ICKDSF – batch job
- ADRDSSU – batch job
- DFSMShsm
- IEFBR14 – batch job
- ISPF 3.*x*.

The required authorizations are as follows:

- Operator reply.
- ALTER access to catalog.
- ALTER access to class DASDVOL (NON-SMS).
- DFSMShsm AUTH userid DATABASEAUTHORITY(USER).
- READ access to class FACILITY profiles.
  - STGADMIN.ADR.STGADMIN.DUMP
  - STGADMIN.ADR.STGADMIN.RESTORE
  - STGADMIN.ADR.STGADMIN.COPY.RENAME
  - STGADMIN.ADR.STGADMIN.xxxx

All the tasks listed above can be performed without the OPERATIONS attribute, provided that

- The FACILITY class is active.
- The RACF profile for the desired function is defined.
- The user has READ access to the profile.

READ access to a STGADMIN.ADR.STGADMIN.** profile and the use of ADRDSSU will grant authorization similar to OPERATIONS, except that it is not possible to read or update datasets. For a complete list of profiles, refer to the chapter entitled

'Protecting DFSMSdss keywords with RACF' in the *DFSMSdss Storage Administration Guide* (SC26-4930-03 or later).

In our shop, we successfully removed all OPERATIONS attributes for TSO users about 10 years ago. So far, we haven't needed to reintroduce any of them.

However, it's important to remember that there might be situations in which a userid with OPERATIONS attribute may be of value (the only situation I can remember is when we were merging two catalogs with IDCAMS REPRO MERGECAT). For this rare occurrence, we keep a surrogate userid which can be used only in a batch job. This userid has not been used for the last three years.

One of the real advantages of not having any TSO OPERATIONS users at all is that all the output from batch jobs can be filed and presented to auditors.

The problem, as so often happens, is to do with people – it's rare to find someone with special authorities who is happy to give them away. Try to start with unnecessary batch OPERATIONS userids, and then convince your auditors. They should support you in granting the lowest level of access that is necessary.

Karl Reinhard Blatt
Systems programmer (Germany)

Looking for a specific article?

If you keep hoping for an article on a particular topic, but we never publish one, please let us know what the subject is. If it's likely to be of interest to other subscribers too, we'll commission it and publish it in *RACF Update*.

Visit the *RACF Update* Web site

http://www.xephon.com/racf

and follow the link to *Opportunities for RACF specialists*.

RACF and security enhancements in new z/OS versions

IBM's announcement of z/OS V1R2 and preview of z/OS V1R3 on 11 September introduced significant security enhancements provided by RACF and other z/OS components. These are outlined below.

RACF-SPECIFIC FUNCTIONS FOR z/OS V1R2

The RACF-specific functions for z/OS V1R2 are as follows:

- The ability to create a new kind of group that can contain an effectively unlimited number of users. This accommodates the need to associate more users under a RACF group definition when designing e-business applications.
- Improved Unix security, through:
  - improved RACF messages for security failures while accessing Unix files and directories.
  - extensions to superuser granularity to cover the chmod command.
- Improved availability, through better toleration of CF errors.
- Improved security tracing, which minimizes the time spent doing problem determination.

RACF-SPECIFIC FUNCTIONS PREVIEWED FOR Z/OS V1R3

The RACF-specific functions previewed for z/OS V1R3 are as follows:

- PKI, a new component of the SecureWay Security Server, will be embedded in z/OS. This consists of:
  - a certificate authority that provides digital credentials to participants.
  - a public-key cryptographic system that uses these digital credentials to help ensure overall message integrity, signature verification, and user authentication.

PKI is generally agreed to be critical for transaction security and integrity. New functions in this release will extend the currently available Web-based front-end to manage the entire life-cycle of a digital certificate that is based on PKI. Using the Web interface, it will be possible to generate digital certificates for both users with RACF user IDs and external clients. Additionally, it will be possible to administer certificates and certificate requests using the same Web-based front end. (Note that this extends the Web-based PKI services already supplied for OS/390 V2R10, z/OS V1R1, and z/OS V1R2.)

- In z/OS V1R3, RACF and Unix will allow the use of access control lists, ACLs, to increase the quality of file and directory access controls by adding extended pemissions assigned to individuals and groups.

The announcement also contains other security-related items, including added SSL, Kerberos, and digital-certificate support, and intrusion detection.

REFERENCES

The full announcement letter can be found at:

```
http://www2.ibmlink.ibm.com/cgi-bin/
master?xh=0iREdwY6y8yOjj1USenGnN9332&request=announcements&parms=H%5f201%2d248&xfr=N
```

The *z/OS V1R2 Overview and Release Guide* can be found at:

```
http://publibz.boulder.ibm.com:80/cgi-bin/bookmgr_OS390/BOOKS/E0Z2A110/
CCONTENTS
```


E-mail alerts

Our e-mail alert service will notify you when new issues of *RACF Update* have been placed on our Web site. If you'd like to sign up, go to http://www.xephon.com/racf and click the 'Receive an e-mail alert' link.

Can RACF help to block the FTP exposure?

FTP is a well-known file transfer application that is available within TCP/IP. Its functionality is defined in an RFC (dated 1985, number 959).

More often than not, security managers do not view FTP as a network facility that could create any security exposure. They are persuaded that it has limited functions and does not give full access to the system. In any case, RACF is supposed to protect datasets and authenticate any entity that wants to initiate a file transfer.

However, as we'll see below, FTP offers more than just the ability to move files from one environment to another.

FTP VERSUS PROPRIETARY PRODUCTS

A decade ago, in the mainframe field, there was no FTP. Instead, there were a variety of proprietary products, which still exist today. In general, these are very effective, and more powerful than FTP because they can do compression, encryption, checkpoint-restart, etc. They offer some security of their own, mailboxes, and interfaces to monitor the transfers, and they are well integrated into current production systems. The drawback is that they are invasive: some part of the product must be installed on any platform involved in the file transfer.

FTP, by contast, is now standard on every system that supports TCP/IP. Many sites want to benefit from FTP because it is a standard tool, which is immediately usable by themselves and their partners at no cost. Some sites even consider dropping the proprietary products they've been using, because these were adapted to the closed world of SNA, and are no longer cost-effective in the IP world. However, the generalized use of FTP raises some security problems (see below).

Now that the mainframe is part of the TCP/IP world, it has become much more accessible than it was in the closed SNA world. There's no doubt that, for the ordinary (or potential) cracker, Telnet is the front door through which to attempt to penetrate the system. However,

many crackers prefer back doors – and FTP may be just the back door that gives undue access to OS/390...

THE NETWORK ASPECT

RFC number 2577, 'FTP Security Considerations', published in 1999, mentions many potential risks stemming from the use of FTP:

- The bounce attack, by which the client instructs the FTP server to send to a machine he wants to attack a file that may contain specific commands (eg SMTP commands).
- Spoofing attacks could defeat protection based on network addresses.
- Brute force password-guessing through the FTP server is possible if you do not limit the number of password attempts.
- Denial of service attacks, the purpose of which is to disable access by a valid user.
- A malicious client may determine valid userids on a server, because the server's response differs when the userid exists and when it does not.
- Passwords being sent in clear text may be subject to eavesdropping.
- More generally, privacy may be at risk, since all data is sent across the network in unencrypted form (standard FTP provides no encryption).
- Port stealing by an attacking client may prevent another legitimate client from making a transfer.

Not all these weaknesses will be considered in this article. For example, the problem of denial of service against userids due to password-guessing is not specific to FTP. Similarly, many of the other risks listed above are common to all TCP/IP-based protocols. Static passwords should be avoided when possible, and many solutions exist today to replace them, including passtickets, tokens, certificates, etc.

Although some security extensions to the FTP protocol were proposed in 1997 (they constitute RFC2228), they are not in widespread use

today. VPNs or FTP over SSL could be the way to go to make FTP much stronger, but this requires more products or facilities to be installed.

FTP ON THE MAINFRAME

From OS/390 2.5 onwards, the IBM TCP/IP stack serves both traditional MVS and USS applications. This means that both MVS and HFS file systems may be accessed through the same FTP port. Before this version, there was a FTP server (port 21) for MVS datasets, and an OE FTP server (often associated with port 1021) to handle OE files.

The alternative to the IBM TCP/IP stack is TCPaccess from CA (originally Interlink).

Although they're very different, the two stacks offer a comparable service. TCPaccess does not require USS. As far as FTP is concerned, TCPaccess controls both inbound and outbound sessions, while with the IBM stack the outbound FTP traffic is not controlled (the FTP client on OS/390 will directly contact the remote FTP server, and the local OS/390 FTP server is not involved).

Although this article discusses my experiences with the IBM stack, most information applies to both stacks.

WHAT DOES RACF OFFER TO BETTER CONTROL FTP?

RACF offers partial protection for FTP. You can control which IP addresses are entitled to access the FTP server. This is done by profiles in class TERMINAL. A profile such as A0* with UACC=NONE will prevent any 160.xxx.yyy.zzz IP address from doing FTP activity (A0 in hexadecimal is 160 in decimal).

For example, a RACF command such as

```
rdefine TERMINAL A0* UACC(NONE) »
```

will prevent any client with address 160.xxx.yyy.zzz from accessing the FTP server:

```
C:\WINDOWS>ftp myhost
Connected to myhost.
220-FTPD1 IBM FTP CS V2R10 at S390, 10:02:29 on 2001-04-29.
220 Connection will close if idle for more than 5 minutes.
User (160.27.47.145:(none)): prod01
331 Send password please.
Password:
530 PASS command failed - __passwd() error : EDC5163I SAF/RACF extract
error.
Login failed.
ftp>
```

On the server side, the following message will pop up in the MVS syslog:

```
ICH408I USER(PROD01  ) GROUP(ETIC) NAME(THIERRY FALISSARD   )
LOGON/JOB INITIATION - NOT AUTHORIZED TO TERMINAL A01B2F91
```

What's more, the date and time can be used to limit access during certain time periods. For example:

```
rdefine   TERMINAL   A01B2F91    WHEN(DAYS(WEEKDAYS)   TIME(0700:1900))
```

However, the protection based on terminals is an all-or-nothing solution, which will apply to any FTP server you may have. You cannot grant to this range of IP addresses access to another FTP server on your system.

In recent versions of RACF, the SERVAUTH resource class enhances TCP/IP security by controlling access to the TCP stack or to TCP ports. You could use it to restrict access to FTP ports to some userids.

Anonymous FTP accesses are possible (thanks to the ANONYMOUS statement in the FTP.DATA configuration). If you allow this, I would recommend you to associate the anonymous clients with a restricted userid, so that they can handle only resources they are explicitly entitled to use, not benefiting from any UACC.

AUDITING

There is no audit facility to instantaneously monitor all FTP-related events. You might undergo attacks such as password-guessing without receiving any notice, with no clue about the originator IP address. Network administrators do not have a clear view of the FTP traffic. The only way to get information is to have SMF records type 118 cut

for FTP servers and FTP client calls. A record can be cut for each FTP command – you have to ask for it in the FTP.DATA dataset. This will enable you to do some after-the-fact auditing.

SOME RISKS WITH FTP

There are a number of risks with FTP:

- A valid userid/password is required to connect to the FTP server on OS/390. On the other hand, RACF access lists actively protect datasets against unauthorized access (this is not true for USS datasets).
- However, security requirements may differ when a userid accesses data from the internal network, or when the same userid is using a client FTP program from the outside or through an insecure network like the Internet. You may want to restrict access to certain files from certain places so that a certain file cannot be transferred out of your organization.
- You may not want confidential data to flow from the mainframe (where it is supposedly secure), through a public network, to the user's home computer. FTP has no encryption feature imbedded into it, and anyway the end user's PC is not the best place for confidential data.
- The user may also try to use FTP to 'navigate' into your catalogs. For example, the CD command enables him to list all datasets with a high-level qualifier of SYS1, and then try to transfer some of them.
- Datasets can be deleted by a FTP command, and, unlike with TSO, no confirmation is asked before deletion.
- A malicious user might transfer big files to the host so as to overuse disk space. When allocation is allowed to everybody, it has been noticed that sooner or later the space serves as a repository place into which hackers deposit files.
- A user may also try repeatedly to mount tape datasets. Too many mount requests may easily disrupt production by delaying jobs:

```
ftp> ls *
200 Port request OK.
125 List started OK
ISPF.ISPPROF
JCL
TEST.TAPE
250 List completed successfully.
ftp> get test.tape
200 Port request OK.
125-Waiting for allocation of tape data set PROD01.TEST.TAPE
```

- Last but not least, many people don't know about one facility that FTP provides on OS/390: job submission. You can submit a JCL that may be stored either on the server or on the client. The way to do this is as follows:

```
ftp> quote site file=jes
200 SITE command was accepted
ftp> get jcl
200 Port request OK.
125-Submitting job jcl FIXrecfm 80
125 When JOB01770 is done, will retrieve its output
250 Transfer completed successfully.
```

  The client gets the job output as soon as the job is terminated (in this example, the job output will be copied to the PC into a file also named 'jcl').

  You may not want to grant this job submission feature – and RACF is of no help here. I remember one site where I had no access to TSO; nevertheless, I was able to do my work through FTP, by submitting jobs and getting back the output!

  We won't examine here a similar facility to issue SQL commands ('quote site file=sql') because some set-up is required for it to work, so only a deliberate decision in your site can allow it.

THE USS ASPECT

Because Unix crackers who try to penetrate through FTP will not necessarily be interested in traditional MVS files, you shouldn't overlook the USS environment. The FTP client simply has to type a CD command (for example CD '/') to get access to it.

In many sites, the USS file system is still a black box; very often, it is less protected than standard OS/390 files. It's not unusual to find Unix

files susceptible to Trojan horse or backdoor attacks, unprotected files with APF authorization (be cautious with facilities like 'Tools and Toys'), files that are program-controlled or have SETUID or SETGID authorization, and many 'goodies' that may come from the open world of Unix.

At the very least, you should review globally writable Unix files, because some may be used by attackers to increase their authority. APF files can potentially subvert the security of your whole system.

FTP EXITS: A GOOD WAY TO PROTECT FTP

What's needed to harden FTP, without diminishing its flexibility, is to be able to check accesses based on IP address, userid/password, FTP port, FTP command, and subcommand. With IBM's FTP, some exits can be used to implement a sort of mainframe-based 'FTP firewall', which is more reassuring for the security manager than an external firewall with fewer controls. These exits are described in turn below:

- The FTCHKIP user exit is called at logon, or whenever a new connection is opened. It can use the IP and PORT addresses of the local and remote hosts to decide whether the remote host's connection should be cancelled.
- The FTCHKCMD user exit is called whenever the user enters a command. It can rely on the user ID, the command, and the command parameters to permit or block the execution of the command. The best control is achieved by combining this with the previous exit (which knows the client's IP address).
- The FTP server SMF user exit, FTPSMFEX, is called before an SMF record that contains information about an FTP server session is written. It could be used to do some real-time auditing, or to log FTP sessions and commands in a specific dataset.

TCPaccess from CA has equivalent exits, namely: FTPCMND to control FTP commands, and FTPLOGIN to control FTP logins.

WHAT ABOUT OUTBOUND TRAFFIC?

Because IBM has implemented only the FTP two-party-model in its

stack, and the FTP client offers no exit, outbound FTP traffic cannot be protected. With TCPaccess, the local FTP server can control outbound FTP.

COMMERCIAL PRODUCTS

You can write your own FTP exits and enforce some security rules to control FTP users. Alternatively, if FTP security is considered critical at your site, you might prefer to buy a dedicated product that would use the FTP exit points to offer the following features:

- Tracing. Every FTP action should be logged in the system log or in a specific file.
- Granularity of protection with dedicated RACF profiles. For example, to forbid job submission from FTP port 21, you might have an FTP.0015.SITE.FILETYPE.JES profile in FACILITY class with uacc=NONE.
- The client's IP address combined with the userid could be a selection criterion. For example:

```
rdefine FACILITY FTP.0015.CWD.* audit(all(READ)) uacc(NONE)
permit FTP.0015.CWD.* class(FACILITY) id(PROD01) access(READ) +
            when(TERMINAL(C0A83211))
```

  This will allow userid PROD01 to issue the CD FTP command only if his IP address is C0A83211 (that is, 192.168.50.17 – generics work too).

I know of two commercial products which can help you do this:

- Secure\FTP from Link\Manage.
- FTPAlert from WDS.

CONCLUSION

The Internet was originally designed to be open. And now we are trying to protect it by closing all the doors. This is a difficult task for us dinosaurs, who are not yet used to working in an open world. It's equally difficult for network administrators with a Unix or NT background, who have to understand the mainframe specifics. But be

warned: FTP is one of these apparently innocuous protocols that actually need appropriate protection, be it through home-grown exits or dedicated products.

REFERENCES

- RFC 959 (File Transfer Protocol), at:

  www.w3.org/Protocols/rfc959/
- RFC2577 (FTP security considerations), at:

  www.faqs.org/rfcs/rfc2577.html
- *Top Ten Ways to Compromise OS/390 Security*, at:

  os390-mvs.hypermart.net/tenflaws.htm
- The SERVAUTH Resource Class, at:

  www.stuhenderson.com/servauth.html
- Secure\FTP – www.linkmanage.be/ps_prod_secureFTP.htm
- FTPAlert – www.willdata.com/htm/ftpalert.htm

Thierry Falissard
(France)

Interested in writing an article, but not sure what on?

We've been asked to commission articles on a variety of RACF-related topics. Visit the *RACF Update* Web site, at:

http://www.xephon.com/racf

and follow the link to *Opportunities for RACF specialists*.

RACF and ACF2 – a comparison

1 **ACF2** can protect dataset members.

RACF can currently protect only on the dataset name.

2 **ACF2** allows testing of access rules.

RACF does not.

3 **ACF2** allows audit trail generation without having to purchase an add-on administrative product.

RACF requires an add-on product (ie Vanguard RACF Administrator or Consul/RACF) to run reports.

4 **ACF2** allows cloning when creating logonids.

RACF does not allow cloning, without the purchase of an additional administrative product (eg VRA or Consul).

5 **RACF** has a global access table, so that rules can be stored. The global access table does not deny access, but will refer to the RACF profile if access is not allowed.

With **ACF2**, dataset rules can be made globally resident, which reduces overhead on validation checks.

6 **ACF2** can restrict the time of day when someone can sign on to the system via the shift field.

RACF cannot.

7 **ACF2** can trace all of an individual's activity (TSO commands or dataset access) as well as log the activity of a resource.

RACF can log resource activity but cannot log individual activity without a written exit.

8 **RACF** with an add-on product (VRA) can determine all the resources to which an individual id has access.

ACF2 does not have this capability.

With VRA, **RACF** can also list all occurrences of an id on access lists that no longer exist.

9 Granting access with **RACF** is as simple as connecting a userid to a group that has access. Users can be connected to multiple access groups.

ACF2 requires a change to the logonid's uidstring or a change to the rule key before access can be modified.

10 **ACF2** maintains a listing of security violations in the logonid record.

RACF does not maintain this field.

11 **ACF2** allows the password violation counter to be set back by one.

RACF does not allow this.

12 **ACF2** rules allow temporary access to a dataset.

RACF does not allow temporary access to a dataset via the profile. What can be done in RACF, to provide temporary access to a resource, is to connect the user to a group that has access to the resource with a specified revoke date.

13 **ACF2** has a READALL privilege that allows browse access to any file (without update/delete).

RACF has only the OPERATIONS privilege, which provides full access to every dataset.

14 **RACF** allows the running of a DSMON utility which captures all of the RACF parameters and lists all RACF Groups and Privileged Userids.

ACF2 has the Global System Options which states all of the ACF2 parameters, but does not list privileged logonids.

15 **ACF2** requires the use of NEXT KEYS when RULE KEYS fill up.

I have not (to date) seen a **RACF** profile run out of space.

Without a NEXT KEY, **ACF2** rules for a High Level Qualifier are all stored under one key.

RACF can have several rule profiles for each High Level Qualifier.

16 **ACF2** and **RACF** now both require passwords for all ids on the system. **ACF2** did not make this requirement previously.

17 **ACF2** protects by default.

RACF does not protect by default, unless PROTECTALL is turned on.

18 **RACF** is reputed to be easier to interface with third-party products than **ACF2**.

19 Control of batch JOB submission is more centralized in **RACF** than in **ACF2**.

RACF controls job submission through one class.

ACF2 controls job submission in several places.

20 **RACF** requires a resource or id to be owned by an id or group.

ACF2 does allow the owner of a resource with the $owner field, but this is optional.

21 **RACF** rules are stored either through the panels or TSO command line, but never from a dataset.

ACF2 rules are usually stored in a dataset, but rules can be compiled directly from ACF2. Because rules can be compiled from a dataset or directly from ACF2, a rule change from one source can wipe out a rule compiled in another source.

In addition, **RACF** rules can be changed by one line command.

ACF2 rule updates require going through a series of commands.

22 **RACF** uses the RVARY command to query or switch the RACF databases. The databases are password protected.

ACF2 has no similar function.

23 **ACF2** allows the use of Firecall Emergency Logonids. When a programmer has an emergency requiring access to production resources, they log on to the firecall id to fix the problem. Their name and id is captured for accountability purposes.

RACF does not have any such provision. An individual can be connected to a group for emergency access, but there is no way to connect someone to a group that has non-cncl (OPERATIONS in RACF) privilege, because a privilege cannot be assigned to a group.

24 Within **RACF**, a user can be connected to several functional groups that have different levels of access.

ACF2 allows masking of the uidstring or lid for rules, but a user cannot be part of multiple groups.

CONCLUSION

If you use RACF alone, without an add-on package (ie Vanguard RACF Administrator or Consul), ACF2 offers more granularity and flexibility. With the addition of VRA or Consul to RACF, the distinction blurs.

In general, ACF2 provides more flexibility and better protection to logonids/userids, while RACF provides better protection to resources.

Bruce Josephs
(USA)

Leaving? You don't have to give up *RACF Update*

You don't have to lose your subscription when you move to another location – let us know your new address, and the name of your successor at your current address, and we will send *RACF Update* to both of you, for the duration of your subscription. There is no charge for the additional copies.

The changing world of IT security

In the first of three articles in this issue looking at remote security, we consider the threats posed by Distributed Denial of Service and virus attacks.

Even if you managed to maintain a purely mainframe computing environment, the day your first 3270 was replaced by a Windows workstation running terminal emulation software was the day RACF stopped providing complete protection.

Earlier this year, the frequency and intensity of Distributed Denial of Service (DDoS) attacks increased markedly. And several viruses/worms have gone from unknown to widespread in as little as 24 hours. In fact, two of the most widespread, SirCAM and Code Red, hit at the same time. The industry press is already calling for a change in approach towards combating viruses/worms and Denial of Service attacks. We have to face the fact that it is no longer practical to combat attacks solely at the ultimate target.

DDOS AT GRC

Very few people are willing to talk about computer security incidents, because they generally represent both an embarrassment to the technical staff involved and a huge public relations nightmare for the organization as a whole.

Steve Gibson of Gibson Research Corporation (grc.com), by contrast, has provided detailed information about the multiple DDoS attacks his site has sustained. The first attack, on 4 May 2000, was traced to a 13 year old boy who, with little knowledge or skill, simply downloaded a few pieces of software from an Internet site for hackers and started using them.

Even though this attack, or at least the target, appears to have been decided on on the spur of the moment, he had been collecting his weapons for some time. His arsenal is a worldwide base of Zombies, Internet-connected PCs that he controls remotely without their rightful

owners' knowledge. In total, 474 Windows-based workstations were used in the first attack on grc.com.

It's not clear whether these machines were originally obtained by the boy gaining access to (hacking) the machine and placing the hacker program file there, or by an e-mail virus that installed the hacker program. But it is known that the hacker program is a single file stored in a Windows directory with a name that's almost indistinguishable from a real Windows component. The program is an Internet Relay Chat (IRC) client that identifies itself on the IRC's Internet site.

The attack itself involves issuing a single command that is simultaneously performed by all available infected machines – perhaps a ping command that sends a large number of very large packets.

Steve tried contacting several of the ISPs whose customers' machines were being used for the attack, but most refused to help. The largest said that they work only with the FBI. When contacted, the FBI were friendly, but said they really needed documented losses in the million dollar range before they could investigate.

TWO MAJOR VIRUSES SIMULTANEOUSLY

The SirCAM virus was discovered on 17 July 2001, and McAfee delivered its weekly virus signature file update that detected it late on 19 July 2001. Within little more than 24 hours, it was widespread among individuals and small businesses. Larger organizations had learned their lessons from previous viruses, most notably the ILOVEYOU virus, and had the necessary processes in place: a firewall that scans incoming e-mail for infected attachments, and a check at least once a day for firewall updates.

This was a good thing, because many of these same organizations were then faced with a particularly nasty worm known as Code Red. Fortunately, it was restricted to a particular piece of Microsoft software (IIS) running on Windows 2000 Server. This meant that it did not affect those likely to be hit by SirCAM: individuals at home and in small businesses without their own Internet infrastructure.

Nearly a week later, when contacted by the media, the FBI stated that they had not opened an investigation into SirCAM because no-one

had come forward with a documented loss in the millions of dollars range. (The Dutch authorities, by contrast, investigated the Anna Kournikova virus that first appeared in February 2001 and then arrested the 20-year-old responsible for it, confiscating all his computer equipment in the process.)

MICROSOFT AND MCAFEE

Both Microsoft and McAfee were criticized about the SirCAM virus. About a year earlier, Microsoft released a poorly conceived security patch for Outlook 2000 that should have prevented SirCAM from infecting workstations that used Outlook to receive e-mail. Office XP's Outlook 2002 was delivered 'out of the box' with substantially the same protection as Outlook 2000 with the security patch.

There are a number of reasons why most people never installed the Outlook 2000 patch. First of all, few people know about the availability of 'updates' to Office. Second, the Web-based Office Update facility –- the easiest way for most people to apply these updates –- was down for several months last winter. Finally, word spread quickly about how the security patch worked. It made it impossible to receive many common types of e-mail attachments. And, once you installed the patch, there was no way to control it, disable it, or uninstall it, short of deleting and reinstalling Office 2000 from scratch. Other software vendors didn't help matters, by continuing to send updates to their software as .exe file attachments to e-mail, which were then blocked by customers' Outlook with the security patch.

What's more, the Outlook security patch was incomplete, in that it didn't block at least one of the executable types (file extensions) of attachments that SirCAM used. Worse still, at least when tested with Office XP's Outlook 2002, Outlook did not display the attachment's correct extension, even on a workstation with Windows 2000 Professional set to always display file extensions.

For all this to make any sense, a little background on SirCAM is required. The e-mail attachment it sends is created by locating a non-executable file on the user's workstation, such as ABSTRACT.TXT, and randomly adding one of five executable extensions: .BAT, .COM,

.EXE, .LNK, or .PIF. The result is a file that appears to have two extensions, such as ABSTRACT.TXT.LNK, and which Outlook displays as ABSTRACT.TXT and even indicates is a 'Text file'.

It took McAfee several days to correct its SirCAM oversight. The default behaviour of its anti-virus products, when asked to perform a full virus scan of a disk drive or directory, had long been to scan only certain file types, significantly reducing the amount of time it took to perform a complete scan. Like Microsoft, SirCAM used file types that McAfee did not scan by default. And, at one point during this time period, McAfee's flagship product VirusScan Online began turning off the Scan All setting that many of us had changed to override the McAfee default.

CALLS FOR CHANGE

By and large, ISPs have failed to follow the lead of organizations with their own Internet infrastructure. Organizations long ago transferred the responsibility for hacker and virus protection from the end user to the network security staff. More recently, the anti-virus software on the desktop has become a fail-safe, only there on the off chance that the firewall lets something through. Individuals, small businesses, and anyone else not running their own mail server(s) and firewall should be behind their ISP's firewall, not unprotected as they are today.

Steve Gibson holds a slightly different view:

- Each individual should prevent the use of his/her workstation in a DDoS attack.
- Microsoft should remove the improved DDoS attack capabilities that are new to Windows XP.
- ISPs should block outgoing DDoS traffic.

Putting all of this in perspective, it is important to remember that the majority of all security breaches still originate internally: carried out by members of your organization. Of course, no one likes to talk about such things, which explains why the numbers vary so widely on just

how it divides up between internal and external threats. Plus, in recent years, it's the external threats that have most visibly disrupted the lives of computer users within organizations.

Jon E Pearkins
(Canada)

Contributing to *RACF Update*

In addition to *RACF Update*, the Xephon family of *Update* publications now includes *CICS Update*, *MVS Update*, *TCP/SNA Update*, *VSAM Update*, *DB2 Update*, *AIX Update*, *Domino Update*, *MQ Update*, *NT Update*, *Oracle Update*, and *TSO/ISPF Update*.

Although the articles published are of a very high standard, the vast majority are not written by professional writers, and we rely heavily on our readers themselves taking the time and trouble to share their experiences with others. Many have discovered that writing an article is not the daunting task that it might appear to be at first glance – and the effort involved is more than offset by our generous terms and conditions.

If you have ever experienced any difficulties with RACF, or made an interesting discovery, you could receive a cash payment, a free subscription to any of our *Updates*, or a credit against any of Xephon's wide range of products and services, simply by telling us all about it.

More information about contributing an article to a Xephon Update, and an explanation of the terms and conditions under which we publish articles, can be found at www.xephon.com/nfc. Alternatively, please write to the editor, Fiona Hewitt, at any of the addresses shown on page 2, or e-mail her at fionah@xephon.com

The anti-virus marketplace – focus on McAfee

RACF does not deal with viruses, but someone has to take responsibility for the protection of each workstation now that terminal emulators (workstation software) have replaced 3270s. That means getting to know at least one anti-virus vendor and its products.

THE ANTI-VIRUS MARKETPLACE

Despite concentrating on organizations rather than individuals, McAfee still ranks number one in worldwide market share among anti-virus software vendors. IDC's July 2001 report gave McAfee 29% of the combined consumer and enterprise market, which totalled $1.4 billion US in 2000. The market grew 25% from 1999, despite a major slowdown in IT spending last year.

In the corporate market, other major players include Computer Associates, Symantec, and Trend Micro. Corporate-only market share figures are hard to come by, but McAfee is somewhere in the 40s, percentage-wise, with the others hard pressed to get into the 20s.

Although all four are large enough to merit consideration, it should be pointed out that corporate size matters for anti-virus software because of the skilled manpower required to act quickly during a virus outbreak. Vendors share a virus information clearing house, but that still leaves the challenge of rapidly responding by updating your anti-virus product to handle the new threat.

WHY MCAFEE?

Some years ago, IBM effectively left the anti-virus business by merging its product and research staff with Symantec, who had purchased Norton Anti-Virus several years earlier. As an existing IBM customer, I managed to get two years of free anti-virus software and signature file updates from Symantec. But installation on the two workstations for which I was responsible did not go well, and I quickly switched to McAfee.

Although McAfee has made a lot of mistakes over the intervening years, it still seems very competent at its core anti-virus work. It may not always be the first vendor with an update to handle the latest virus, but it seems much more consistent in the sense of never being too late (as at least one of the other major players has been).

Recently, McAfee began a concerted attempt to move customers off its VirusScan traditional desktop-based software product to several incarnations of its Web-based 'service'. Initially, anti-virus protection was bundled in Clinic, but it can now be purchased separately as VirusScan Online.

KEEPING UP TO DATE

As long as I've used it, McAfee has offered free virus signature updates for VirusScan on its Web site as a download. Updated software had to be downloaded periodically from my telco's intranet, or the latest copy of the product purchased from a retailer.

McAfee then began to offer what it called SUPERDAT updates. SUPERDAT is an executable program you download from its Web site. When you run it, it updates both the software and the virus signature files. Actually, it only updates what McAfee refers to as the 'engine'. Every year or two, McAfee brings out a new version of VirusScan and not all SUPERDAT engine updates are available for the previous version(s), eventually forcing a reinstall from your employer or a retail purchase of the new version.

But all of this is a manual process, at least in terms of remembering to download updates on a regular basis. And, while this may not have mattered in the mid-1990s, when most of us were happy receiving and applying quarterly virus updates from IBM, it's certainly not enough nowadays, when we're more likely to ask "is once a day enough?".

In late 1999, McAfee addressed this problem by licensing BackWeb technology, which ran constantly on your workstation, monitoring for McAfee updates on a minute by minute basis. The download was performed in the background. When it was complete, the BackWeb icon in the system tray alerted you that something needed your attention, and you could then initiate installation whenever it was convenient.

Unfortunately, McAfee discontinued its use of BackWeb just before releasing Clinic, its Web-based service designed to replace the traditional shrink-wrap VirusScan software. With Clinic and its VirusScan Online component, you choose a time each day when an automatic check for updates will be made.

VIRUSSCAN ONLINE

Like other products, VirusScan Online includes a monitor and a scan module. The monitor's main task is to run whenever a workstation file is about to be opened, whether it's an application requesting the open or the operating system itself. It checks the file for virus infection before allowing the open to proceed. If the file is infected, the user is prompted to determine what action should be taken: delete, rename, disinfect, etc. The scan module is run on demand by the user requesting a file, directory, disk drive, diskette, or data CD be checked for virus-infected files.

VirusScan Online is different from traditional software in that it downloads and installs directly. When you click Scan in the McAfee.com Services menu on your desktop, a Web browser window is initiated. You are automatically logged on, to ensure you are licensed to use the 'service', then a check is made to see if the latest scan module, engine and virus signatures are installed on the workstation. If not, they are downloaded and installed. Then the scan module runs in a Web browser window with a GUI that lets you choose what you want to scan. You can even select options and have them remembered for future use. The status of the scan, right down to the name of the current file being scanned, is shown right within the Web browser window.

The monitor updates are a little less well thought out. The daily check opens up several Web browser windows that can disrupt your work even if there are no new updates available. When there is an update available, you are given 60 seconds to click on it, or it downloads, installs, and prompts for a reboot automatically. Over the past year, this update procedure has changed repeatedly, sometimes eliminating the need to reboot, but mostly requiring it.

If your workstation is not running at the time you've selected for a

daily update, it may or may not occur when the workstation is next activated. (Admittedly, my testing has involved the use of the Windows 2000 Hibernate feature, which is probably not used enough for McAfee to have considered it.) Even more baffling is the fact that the update check occasionally initiates a display of a McAfee marketing Web page, complete with pop-up window, rather than checking for an update. On several occasions, an update that failed somehow was considered done, and I had to wait a week for the next update to be released before I could get the previous week's update (updates are cumulative).

Editor's note: 'Virus protection – an essential piece of the mainframe security puzzle' on pages 45-59 of this issue looks in detail at how to implement the e50 as part of the corporate security solution.

Jon E Pearkins
(Canada)

Virus protection – an essential piece of the mainframe security puzzle

Our series of articles on remote security has looked at the problems faced by the small remote office or single user at home trying to securely access the corporate network and RACF-protected mainframe. Here, we consider how the McAfee e50 anti-virus solution can help protect both individual users, and, by extension, the corporate mainframe.

THE PROBLEM

First, one question needs to be addressed: aren't viruses irrelevant in the mainframe world? After all, as IBM likes to brag, there has never been a virus on the mainframe.

Virus-protecting the workstations that run 3270 terminal emulation software to access the mainframe is essential to mainframe security. Why? Because infecting a workstation is the first step that hackers use to gain access to a workstation from the Internet. Today, the headlines and hacker focus is on remotely using other people's workstations for Distributed Denial of Service (DDoS) attacks, but the same approach can be used to gain access to your mainframe from that same workstation. And RACF will not help you if the workstation is already logged on when the user gets called away unexpectedly and the hacker sees his opportunity.

Here is how it could work for a remote workstation with commercial high-speed Internet and a Virtual Private Network (VPN) connection across the Internet to the corporate mainframe. A virus enters the workstation via the user's personal or corporate e-mail account, or software downloaded from the Internet or borrowed from a friend on CD-ROM or diskette. The virus installs a program that makes its presence known to a hacker through Internet Relay Chat (IRC). The same program also accepts commands from the hacker, who can use the commands to monitor the workstation until the user leaves the workstation unattended but still logged on to the mainframe. At this point, the hacker can do anything the user can, from accessing confidential information to (more likely) deleting large quantities of

data. In my first hacker encounter (1980), a 14-year-old 2,000 miles away was forcibly logged off just as he tried to zero the disk drive housing half the corporate data.

BACKGROUND

Last issue, we looked at the WatchGuard SOHO firewall as a way to protect the remote office or single home user accessing the corporate network and RACF-protected mainframe via an ISP's standard high-speed Internet. As both the vendor and the article tried to emphasize, the hardware firewall's focus was illegal access, typically hackers. Not viruses.

When McAfee recently began a strong marketing push for its new line of hardware anti-virus solutions, the parallel was inescapable. Could the low-end McAfee e50 be paired with a firewall to provide complete protection for a remote user or user group? And be ignored once it had been installed?

At the moment, the answer is "No", but that doesn't mean the device should be ignored. It holds a lot of promise both now, in other environments, and in the future as a big step towards a remote security solution. For the moment, however, the e50 really does need to be in an environment where it and the mail server(s) are on the same side of a firewall.

WHAT YOU GET

The e50 arrives in a 9.5" x 17.5" x 20.5" shipping box weighing 28 lb. Open the box and the contents weigh 22 lb, of which the case weighs 15 lb. The case is smaller than a normal PC and is not rack-mountable without additional rack mounting hardware.

A small box contains a mini-keyboard with built-in touchpad to eliminate the need for a mouse, which plugs into the e50's keyboard and mouse ports. Although it weighs about the same as a typical keyboard, it is about the size of a sheet of paper (11.25" x 8.5") rather than the 17.25" x 6.75" of a keyboard, plus additional space for a mouse.

You supply the monitor. If this is starting to sound like a PC, you're right. It is.

Without opening the case, a diskette drive and CD-ROM drive are visible on the front. On the back are two NIC connectors, labelled CARD[1] and CARD[2], PS/2 mouse and keyboard ports, three analogue jacks from a sound card, two USB ports, a parallel port, a serial port, a monitor port, and a red sticker covering the main ventilation (back of the power supply) “Warning, this unit is set for 115 v”.

There are two screws on the back for removing the cover from the case. Inside you’ll find:

- Intel D815BN motherboard with built-in NIC, video and sound cards, and serial, parallel, primary, and secondary IDE, and diskette ports.
- Pentium 633MHz Celeron processor.
- 256MB RAM.
- 20.4GB 7200 rpm Seagate Barracuda Model ST320420A Ultra ATA (66MHz) hard drive.
- Another 20GB hard drive hidden below the Seagate.

There are also two 10-foot RJ45 CAT5 network cables. And four power cords, for four of the international power standards. The manuals and other printed material include:

- *Windows NT 4.0 Workstation Basics and Installation* manual with a Microsoft Certificate of Authenticity with Product ID on the front cover (118 pages).
- *WebShield e50 Installation Guide* (119 pages).
- *WebShield SMTP Administrator’s Guide* (202 pages).
- *WebShield e50 Administrator’s Toolbox* CD-ROM.
- *Read This First* card covering installation and configuration, installing the appliance, and configuring the appliance.
- *Release Notes* for McAfee WebShield e50.
- Sealed PrimeSupport envelope with certificate inside for connect service which provides unlimited toll-free telephone access to technical support during business hours (12 hours per day).

- Several other license agreements, a correction to a manual, and PrimeSupport power and safety information.

DECIDING HOW TO USE IT

Beginning with the Read This First card, you are directed to page 20 of the WebShield e50 Installation Guide, 'Getting started with the appliance'. After stating "You can use the appliance in almost any SMTP network topology", it describes five common situations, all of which have the e50 right beside the mail server, protected from the Internet by a firewall. The only exception is a portion of the fifth topology, labelled an international organization, with three e50s in three different locations.

One warning from the manual is worth noting: "McAfee strongly recommends that you use the WebShield e50 appliance inside your organization – behind a correctly configured firewall – for security reasons."

GATHERING CONFIGURATION INFORMATION

The next recommended step is to determine all names, addresses, and other information needed during installation, to save major delays during the actual installation process. First is the computer name for the e50. The default is E50, but the manual recommends changing it to reduce the likelihood of hackers targeting it. I selected 'MAILWALL'.

Next is network addresses, with this caveat: "McAfee recommends that you do not choose to obtain the addresses from a DHCP server, as this can cause configuration errors in your e-mail set-up. Using static addresses avoids such errors."

There are two NICs in the e50, to support a multi-homed network: two different networks. The manual makes it clear that only one NIC should be used in a single-homed network such as the one I was testing, so only the TCP/IP addresses listed for Network adapter 1 were relevant, namely:

- IP address
- Subnet mask

- Default router/gateway (firewall)
- DNS address for Internet (external)
- DNS address for Intranet (internal)
- WINS server address.

The last three were listed as optional, but with no additional information. Fortunately, the information is readily available by simply plugging a Windows NT/2000/XP workstation into the firewall port where you plan to connect the e50, letting it connect, and then opening up a Command Prompt (Start button-Programs-Accessories-Command Prompt) and typing:

```
ipconfig /all
```

In Windows 95/98/Me, winipcfg replaces ipconfig.

In Windows 2000 Professional, you'll see the following sort of information:

```
Microsoft Windows 2000 [Version 5.00.2195]
(C) Copyright 1985-2000 Microsoft Corp.

C:\>ipconfig /all

Windows 2000 IP Configuration

 Host Name . . . . . . . . . . . . : adiant
 Primary DNS Suffix  . . . . . . . :
 Node Type . . . . . . . . . . . . : Broadcast
 IP Routing Enabled. . . . . . . . : No
 WINS Proxy Enabled. . . . . . . . : No

Ethernet adapter Local Area Connection 3:

 Connection-specific DNS Suffix:
 Description . . . . . . . . . . . : SOHOware 10/100 PCI Network Adapter
 Physical Address. . . . . . . . . : 00-80-C6-E8-33-89
 DHCP Enabled. . . . . . . . . . . : Yes
 Autoconfiguration Enabled . . . . : Yes
 IP Address. . . . . . . . . . . . : 192.168.111.2
 Subnet Mask . . . . . . . . . . . : 255.255.255.0
 Default Gateway . . . . . . . . . : 192.168.111.1
 DHCP Server . . . . . . . . . . . : 192.168.111.1
 DNS Servers . . . . . . . . . . . : 199.185.220.36
                                     199.185.220.52
```

```
 Lease Obtained. . . . . : Tuesday, September 18, 2001 8:20:46AM
 Lease Expires . . . . . : Wednesday, September 19, 2001 8:20:46AM

C:\>
```

I soon learnt that I couldn't just pretend that my firewall used static IP addresses, as the e50 manual suggests. Even though I could reliably predict what IP address would be assigned by my firewall to its internal hub's ports, it was still set up to use DHCP and would not fully activate a port without a DHCP request.

E-MAIL DELIVERY METHOD

Next is the selection of the e-mail delivery method:

- Forward all e-mail messages to a single mail server or gateway.
- Relay all e-mail messages to a number of different mail servers or gateways.
- Resolve where e-mail messages should be sent using DNS.

I found this confusing, because my ISP has the mail servers that I use, one for incoming (POP3) and one for outgoing (SMTP) e-mail. A third is provided for IMAP.

The first choice (forward) allowed only one host name or IP address to be specified. The second (relay) didn't seem to allow a separation between ingoing and outgoing e-mail, but between domain names.

THE PHYSICAL INSTALL

The *Installation Guide* warns you not to place a monitor on top of the e50 case. For safety, I also removed the power warning sticker because of my concern for adequate cooling. The manual also emphasizes the importance of connecting the LAN cable to one of the NICs, labelled Card[1], and not the other, in a single-homed network.

Installation Step 10, towards the end, has a note marked IMPORTANT, which states, among other things: "Ensure that the appliance's CD-ROM and 3.5" disk drives are empty."

Having tried, and failed, to open the drives manually with a small

screwdriver, I switched on the appliance just long enough to get the tray to open. It was empty.

POWERING UP

When you power up the e50, several screens appear automatically in succession, followed by the familiar 'Microsoft Windows NT Workstation 4.0 with Microsoft Internet Explorer' blue and black graphics window in the middle of the screen. It has a 'Begin Logon' dialogue box over the top of it with 'Press Ctrl + Alt + Delete to log on' and the animated graphic of a hand moving towards three keys on a keyboard (which took some getting used to).

A dialogue box then appeared, labelled McAfee WebShield SMTP e50, listing the ID (Administrator) and password. The message also indicates that the all-upper-case password is case-sensitive, and suggests the password be changed for security reasons. If you wait more than a couple of minutes, without clicking the OK button, you get back to the Logon dialogue box, where Administrator is already filled in. All you have to do is type the password. If you have only ever used a mouse, it also takes some getting used to the touchpad and buttons below it. (Don't forget that hitting Enter on any Windows keyboard is the same as single-mouse-clicking on the default dialogue box button.)

CONFIGURATION WIZARD

The NT desktop appears with, amongst other things, a dialogue box labelled 'Configuration Method', which allows you to choose between configuring manually and automatically (manually is the default).

Another small dialogue box appears, with the message 'Please configure your IP and subnet mask addresses'. When I hit the OK button a larger Server configuration dialogue box was displayed, with a photograph of the back of the e50, with the two NIC ports circled, asking 'How do you want to configure your WebShield e50?'. Single homed was the default choice; multi homed the other. For single homed, I was again reminded to connect only to the NIC port labelled Card[1].

Hitting OK revealed the familiar Network dialogue box of NT 4.0 with the Protocols tab selected and TCP/IP Protocol highlighted as the sole entry under Network Protocols. When, as instructed in the manual, I hit the Properties... button, the Microsoft TCP/IP Properties dialogue box was displayed with the IP Address tab showing. As the manual indicated, the NIC labelled Card[2] was displayed in the Adapter box:

```
[2]Intel 82559 Fast Ethernet LAN on Motherboard
```

Selecting the first entry from the drop-down list displayed the parameters for Card[1]:

```
[1]Intel PRO/100+ Management Adapter with Alert on LAN
```

CHANGING THE COMPUTER NAME

My next task was to change the computer name (not required, but recommended). I hit the OK button to return to the Network dialogue box, and then clicked on the Identification tab. In the Computer Name field, I entered my chosen MAILWALL name, replacing the default E50. It's easy to forget, if you haven't used these dialogue boxes lately in NT 4.0 and don't read the e50 installation manual carefully, that you need to hit the Change button before trying to enter a new computer name. Unlike most read-only fields, the cursor appears when you click on the field, giving you the false impression that you can directly enter values into the field.

I was then prompted to shut down and restart my computer before the new settings would take effect.

AFTER REBOOT

The final step of the 'Manual Configuration' section of the e50 *Installation Guide* directs you to the 'Configuring the WebShield SMTP settings' section of the manual where additional planning is required.

After detailed reading of this section and other sections referenced, including quite a bit in the *WebShield SMTP Administrator's Guide*, I decided that the approach of forwarding all e-mail on to a single Mail Server/Gateway would probably work. Although it referred to both

incoming and outgoing e-mail, it looked as if, by specifying my ISP's outbound SMTP mail server, I would be able to initially run the e50 just on outgoing e-mail from a test workstation with Outlook 2000 set up to direct its e-mail to the e50.

HOW THE E50 IS SET UP

Meanwhile, checking the e50 itself, My Computer on the NT Desktop indicated that C: was a 7.84GB partition labelled SYSTEM with 561MB in use, and D: was a 9.76GB partition labelled DATA with 4.74MB in use. Both were NTFS. NT Disk Administrator showed two hard drives, each 19454MB, but each with one partition that was significantly smaller than the drive capacity.

Another point to note is that there is some sort of blanking screen saver in place that makes you think the e50 may have failed if you return after an extended period and see a blank monitor. The monitor is not in the normal sleep state you expect to see when Power Management initiates it on a modern Windows workstation.

After all my detailed reading, I went back to the 'Please configure WebShield e50' dialogue box, hit the OK button, and the WebShield SMTP Configuration Wizard appeared.

SELECTING E-MAIL FORWARDING

When I selected the forwarding option, I was prompted to enter the name or IP address of the server to relay all e-mail messages to. But when I did this, a server diagnostics dialogue box appeared with the following message:

```
Finding Server:
smtp.telusplanet.net
Cannot resolve the given machine name.
```

There was only an OK button, and no diagnostic details were provided in the e50 *Installation Guide*.

I decided to check whether the e50 had been able to connect properly to the Internet through the firewall. I'd seen a reference to Internet Explorer (IE) Version 2 in the Release Notes, but there was no sign of IE in the Start button menu or on the desktop. Start-Programs-

Command Prompt and an ftp to a known site gave me an 'Unknown host' message, and a ping to a known site gave a 'Bad IP address' message. Yet an ipconfig /all indicated NT had assigned the expected IP address, and double-clicking on Network Neighbourhood on the desktop, then Entire Network, Microsoft Windows Network and Workgroup, successfully found the other workstation connected to the same firewall.

I decided to focus on getting a working connection to the Internet. To remove one layer of complexity, I decided to bypass the firewall and connect directly to the ISP by connecting the e50 NIC to the ADSL modem using a CAT5 crossover cable.

DHCP, NOT STATIC IP

I quickly realized that this would mean going back to the Control Panel settings again, but this time without a wizard to guide me. Then it finally sank in that the firewall was set up for DHCP and couldn't be expected to work with a workstation set up for static IP addresses.

After reconnecting the cables, I pushed the Start button, then Settings-Control Panel-Network, Protocols tab, and the Properties button. The Microsoft TCP/IP Properties dialogue box appeared with the IP Address tab displayed. I selected Adapter 1 from the drop-down list under 'Adapter', then 'Obtain an IP address from a DHCP server'. A Microsoft TCP/IP dialogue box immediately appeared with the message:

```
DHCP protocol will attempt to automatically configure your workstation
during system initialization. Any parameters specified in these property
pages will override any values obtained by DHCP. Do you want to enable
DHCP?
```

Once again, I was prompted to shut down and restart my computer before the new settings would take effect.

REBOOT AND IT WORKS

The e50 rebooted as it should, without any manual assistance, and this time my ISP's outgoing SMTP mail server passed inspection.

The next dialogue box to appear read:

```
Webshield can be configured remotely by setting up trusted clients of
the Webshield configuration service.
Specify the trusted clients able to remotely configure this WebShield
Server:
```

There was already a single entry, localhost, and the *Installation Guide* stated: “The ‘localhost’ entry identifies the appliance, and must not be removed from the list.”

Hitting the Next> button displayed another wizard dialogue box announcing that Webshield could send notification of infected e-mail messages to a designated e-mail address, and allowing me to specify the address required (typically the postmaster of your domain). Although this seemed a very useful feature, I left the box unchecked to simplify things initially.

The last dialogue box of the wizard stated:

```
This concludes the WebShield SMTP Wizard. Press the 'Finish' button to
configure WebShield with the choices you have made, or 'Cancel' to
abandon the changes. If you abandon changes, you can manually configure
WebShield using the WebShield Console later.
```

Although this last sentence may be technically correct, it fails to mention that hitting Cancel will take you back to the ‘Please configure WebShield e50’ dialogue box, with the only way out, other than shutting down the e50, the OK button which starts up the wizard.

There were actually three buttons on this last wizard dialogue box: Back, Finish, and Cancel. I selected Finish.

CONFIGURATION COMPLETE

The WebShield e50 dialogue box then reappeared, saying that WebShield e50 configured successfully. The larger WebShield e50 Installer dialogue box underneath it had much more to say:

```
WebShield e5Ø installation process has started.
NT Network has been configured, but WebShield SMTP has not.
Mailscan service startup set to 'automatic'.
Mailconfig server startup set to 'automatic'.
Outbreak Manager service startup set to 'automatic'.
Network configuration complete, WebShield SMTP configuration complete.
```

At the same time, a new logo joined the NIC card logo and the time

in the system tray in the lower right corner of the screen. The new logo was a dark blue circle with a red centre.

UPDATING VIRUS SIGNATURES

Hitting the OK button on the small dialogue box displayed a larger WebShield e50 dialogue box urging me to configure autoupdate/ upgrade, to enable me to remain up-to-date with the latest virus definitions.

I hit yes, to fully reveal the WebShield Configuration Console. The Server was selected in the left sidebar, the first of eight items, referred to in the *Installation Guide* as 'modules': server, delivery, scanning, exclusion, anti spam-relay, blocking, content filtering, and logging. Anti spam-relay and blocking had red lines through them.

On the right there were four tabs: information, automatic update, trusted clients, and service monitor. When I clicked on the automatic update, I was given two choices under transfer method:

- Copy from a local network computer
- Download from the Internet (the default).

I selected 'download' and pushed the 'configure' button. After a couple of other boxes, a Scheduled Update Configuration dialogue box appeared with four options, all, by default, empty check boxes:

- Check for updates at system start-up.
- Use scheduler.
- Randomize updates to within one hour.
- Update silently.

I selected all but randomize, which is intended to avoid network contention when multiple e50s would otherwise be updating at the same time. The 'use scheduler' option includes an area labelled 'How often would you like to check for updates?', with defaults of every 7 days and at 9 am.

HOURLY, NOT WEEKLY

I opted to check hourly, and assumed, since I could find no more information in any of the Guides, that the time field was now irrelevant. However, when the Upgrade Selections dialogue box appeared, completing the wizard and stating that updates were scheduled for every 1 hour starting from 9am, alarm bells began to ring. To be on the safe side, I changed the time to 1 am. I was then prompted to restart the WebShield SMTP MailCfg service on localhost and the WebShield SMTP Service.

NO MORE WIZARDS

This was the first time that the installation process had not been controlled by a wizard, and a bare desktop was visible. As well as the usual icons on my desktop, I now had icons for back-up and restore, configuration console, status monitor, and install Adobe Acrobat reader.

But the majority of the desktop was occupied by "McAfee" in large lettering, "THE VACCINE FOR E-BUSINESS" just below it, a three-shield logo, and "WebShield e50" just below the logo.

The next installation step listed in the manual was modifying the Domain Name Server (DNS), which is required only if a DNS determines where messages should be sent. The final step was testing the configuration. Two tests are detailed in the e50 *Installation Guide*:

- Verifying that the appliance's MailScan service is running.
- Verifying that e-mail messages are scanned and delivered correctly.

TESTING THE CONFIGURATION

The first test involved using telnet to access the SMTP port of the e50 from another workstation. Since DHCP was used, this requires running ipconfig /all on the e50 to determine the IP address. The relevant part reads:

```
Ethernet adapter E100B1:

   Description . . . . . : Intel(R) PRO Adapter
   Physical Address. . . . . : 00-D0-B7-49-D8-1E
```

```
DHCP Enabled. . . . . : Yes
IP Address. . . . . : 192.168.111.4
Subnet Mask . . . . . : 255.255.255.0
Default Gateway . . . . . : 192.168.111.1
DHCP Server . . . . . : 192.168.111.1
Lease Obtained. . . . . : Sunday, September 16, 2001 2:01:10 PM
Lease Expires . . . . . : Monday, September 17, 2001 2:01:10 PM
```

At a Command Prompt on a Windows 2000 Professional workstation running off another port on the same firewall, typing

```
telnet 192.168.111.2 25
```

clears the screen, then, after about 15 seconds, displays one long line:

```
220 mailwall WebShield SMTP V4.5 Network Associates, Inc. Ready at Sun
Sep 16 21:49:04 2001
```

This matches the expected response listed in the *Installation Guide*.

The second test involves creating an EICAR test file and sending it as an e-mail attachment from a workstation to see if it is received and handled by the e50. EICAR is the European Institute of Computer Anti-virus Research, a coalition of anti-virus vendors headquartered in Europe. The test file is a single line in a text file. According to the manual, it's treated like a virus by all anti-virus programs. The file must be named EICAR.COM and the contents must be a single line consisting of the following:

```
X50!P%@AP[4\PZX54(P^)7CC)7)$EICAR-STANDARD-ANTIVIRUS-TEST-FILE!$H+H*
```

However, when I sent a test e-mail with this file as an attachment, it arrived safely at its destination, complete with attachment. And McAfee VirusScan Online didn't indicate that the attachment was infected, so there may be a typo in the e50 manual's listing of the contents of the EICAR.COM file. Nonetheless, the test did indicate that the configuration was correct for outgoing mail.

For a summary view, the WebShield Status Monitor has a Statistics section that provides counts of: received, scanned, filtered, delivered, infected, corrupted, deferred, cleaned, blocked, returned, and quarantined. Mail logs can also be browsed from the Status Monitor.

CONCLUSIONS

As we saw earlier, virus protection is essential for workstations

accessing the mainframe and for mainframe security itself, because infecting a workstation is the first step that hackers would use to remotely gain access to the mainframe.

Although the e50 looks like a promising solution for larger environments, with their own on-site mail server(s), it isn't really suitable for the small remote user(s) office connecting to the RACF-protected mainframe.

There is also the issue of price. The e50 provides protection for up to 100 users for $3640/£2480, including the first year of support. If you're looking at it for a single user or very small group, the per user cost may well be too high.

One final and important point that is easy to miss when evaluating the e50: although the update process, as McAfee likes to call it, very effectively automatically updates the virus signature files as frequently as hourly, the upgrade process, where the anti-virus engine is updated, must be initiated manually.

Jon E Pearkins
(Canada)

Call for papers – share your expertise and earn money at the same time!

Why not share your expertise and earn money at the same time? *RACF Update* is looking for technical articles on mainframe security issues and developments and sample code that experienced RACF practitioners have written to make their life, or the lives of their users, easier.

Articles can be of any length and can be sent or e-mailed to Fiona Hewitt at any of the addresses shown on page 2.

More information about how to contribute can be obtained from our Web site, at www.xephon.com/nfc

IBM's announcement of z/OS V1R2 and preview of z/OS V1R3 introduced significant security enhancements provided by RACF and other z/OS components.

For further details, see the article entitled 'z/OS RACF enhancements' on pages 21-22 of this issue.

* * *

Tivoli Intrusion Manager is a new entry-level security product that detects potential vulnerabilities and provides a central event and problem management console for monitoring intrusions on up to 20 systems.

It uses DB2 Universal Database (UDB) and runs on Windows 2000 Professional, Server and Advanced Server, and NT 4.0. Version 3.7 is the first release of the product.

For more information, contact:
Tivoli Systems, 9442 Capital of Texas Highway North, Arboretum Plaza One, Austin, Texas 78759, USA.
Tel: (512) 436 8000.
URL: http://www.tivoli.com

* * *

eTrust Internet Access offers modular Web access security that authenticates users connecting over the Web and controls their access to business applications and data.

eTrust Defense combines virus protection, danger mobile code prevention, firewall protection, and intrusion detection for Internet gateways, messaging systems, enterprise servers, and user desktops.

eTrust Management is an integrated bundling of security products for large organizations: eTrust Audit, eTrust Policy Compliance, eTrust CA-ACF2, eTrust CA-Top Secret, eTrust Single Sign-On, eTrust Admin, and eTrust Access Control.

For more information, contact:
CA, One Computer Associates Plaza, Islandia, NY 11749, USA.
Tel: (800) 225 5224.
URL: http://www.ca.com

* * *