

INTERDEPARTMENTAL

MASSACHUSETTS INSTITUTE OF TECHNOLOGY CAMBRIDGE, MASS. 02139

from the office of

To: R.L. Bisbey II
C.T. Clingen
F.J. Corbató
F.C. Daley
J.W. Gintell
B.S. Greenberg
D.M. Jordan
R.A. Roach
M.D. Schroeder
J.C. Whitmore
T.H. VanVleck
D.R. Vinograd

From: J.H. Saltzer

Date: February 9, 1973 (Revision 5)

Subject: List of Multics Security Holes

Enclosed is a list of all outstanding security holes in Multics which have been reported to me, and for which fixes have not yet been installed on the current service (645) system. For many of these problems, fixes have already been coded, and are either awaiting installation, or are to be installed only on the 6180 system.

It is interesting that so far, no problems have been found which require any rethinking of system organization. Also, no problem yet found permits a systematic attacker to read a given target segment.

I would appreciate it if the installation and assurance crew would let me know whenever they install a system which they believe fixes a problem documented here, so that I may remove it from the list.

Since Multics is used in production at M.I.T. and elsewhere, it seems wise to regard this list as sensitive; please do not pass it on to anyone else without letting me know.

Outstanding security problems in current Multics implementation

<u>problem name</u>	<u>nature of problem</u>	<u>fix</u>
ITT overflow	Repeated calls to hcs_\$wakeup can cause messages to pile up in the ITT, eventually filling it and causing system messages to be lost, and a system crash, denying service to legitimate users.	<ol style="list-style-type: none">1. (Quick) Place a limit on the number of ITT entries which may be queued for a single process.2. (long range) Transmit user-originated messages directly from one user to another, eliminating need for ITT.
Reused address	Following a system crash, salvager may find a single disk or drum page being used by two different page tables. It awards the page to the first page table which it noticed, possibly exposing information to an unauthorized user.	When a reused disk or drum address is discovered, reset the page contents to zero, and delete it from second (and any later) page tables.
Accounting system hierarchy scan	The accounting system, in order to charge for secondary storage use, must read every system directory, since page-second storage counts are stored in the directories. Thus all system directories and files must be made accessible to the operator of the accounting system.	Revise directories to contain account numbers, and store usage information in separate accounting files, which can be accessible to the operator of the accounting system.
Operator login window	When bootloading Multics, the operator must dial a telephone number to login the initializer console. A hostile user could call this telephone number and take over the system as it comes up.	Although a routine practice of dialing in the initializer before bootstrapping could prevent this problem, it would be safer for the system to display a uniquely generated 5-letter password on the operations console, which must be typed in at the initializer to authenticate the operator's identity.
Absentee/Daemon overload	Any user may submit an arbitrary number of absentee jobs or I/O daemon requests, effectively denying absentee or I/O daemon service to authorized users.	Place an administrative limit on number of outstanding absentee jobs and I/O daemon requests that a single user is allowed.

<u>problem name</u>	<u>nature of problem</u>	<u>fix</u>
Metering gates expose everything	The entry point used for reading out system performance meters allows readout of any ring zero program or data base including, for example, typewriter input buffers, which may contain typed passwords.	Need a separate entry point which reads out only legitimate system performance meters. General peek gate then does not require such a long list of persons authorized to use it.
IPC event channel loop bug	If an IPC message with an illegal event channel name is sent to another process, the other process will loop on out of bounds error inside IPC. If sent to initializer, will cripple system.	Add bounds check to IPC message handling procedure.
fsdct update problem	If a device becomes completely assigned, a corresponding page of the fsdct becomes completely filled with zeros. The procedure to update the drum copy of the fsdct notices the page of zeros and discards the page rather than writing it out. The next reference to the fsdct crashes the system. A user with a very large storage allotment could in principle cause this bug to occur.	Provide entry to page copy procedure to force copying, even if contents are zero. (Alternatively, provide code to recover from missing fsdct pages by supplying a page of zeros.)
login table overflow	The list of logins during a single boot-load of Multics is stored in a single segment with no overflow procedure. A single user, by logging in several thousand times, can overflow the segment, making further logins by authorized users impossible.	<ol style="list-style-type: none"> 1. (Quick) schedule frequent system shut-downs and reboot loads. 2. (long range) provide accounting with an overflow procedure to go on to another segment.
Magic numbers in page control	An old hardware bug trap places magic numbers in core where a page is to be read in, then after reading the page checks the numbers. If still there, it assumes the page didn't come in, and reports a page read error to the user. If a user places contrived names containing the magic bit patterns strategically in a directory to which he has only append access, he can effectively delete other entries in the directory.	Remove old trap. (Already can be disabled under control of a CPU switch.)

<u>problem name</u>	<u>nature of problem</u>	<u>fix</u>
Retriever acl-setting bug	The retriever, if it has any difficulty reading an access control list, uses a default list of rewa *.*.* without warning.	Retriever should use default of no access, and refer owner of the segment to a system administrator who can re-adjust access to correct setting.
High speed line carrier detect problem	When in output mode, the system cannot detect that the user has hung up his 1200 baud line, since there is no carrier detect feature on the 202C6 dataset. Another user can then dial in and continue to use the line, with access to previous users' files.	Obtain datasets providing a carrier detect feature, and add software to log out user on carrier failure, just as for low speed typewriter lines.
Linker bug	Certain types of incorrect link definitions will cause linker to go into a loop inside ring zero.	<ol style="list-style-type: none"> 1. Fix linker to accept only valid definitions. 2. Add time out in ring zero to catch all such problems. 3. (On followon) Move linker out of ring zero.
Mailbox is open bug	Current mail command implementation requires that permission to send mail to a user must be coupled with permission to read and delete any mail in the users mailbox.	Revise mail command to use message segment rather than a directly writeable mailbox.
process directory record quota overflow bug	When process directory exceeds record quota, signaller uses wrong stack, crashing system.	fix stack switching bug.
ect terminate bug	The event channel table is created by IPC without properly setting ring numbers. As a result, a user can terminate the ect, get some other (ring-0 accessible) segment initiated in its place, and cause IPC to write in this other segments.	fix IPC to correctly set ACL on ECT.

problem name

process_id argument
validation

nature of problem

the low-order bits of a process identifier are actually the offset of that process' entry in the Active Process Table. Entry point hcs_\$wakeup does not verify that the process_id given as an argument is a legitimate value for a table offset. It does look for the process id at that location, but this check is not foolproof.

fix

short-term: Check offset value to see if it is legal.

long-term: Use a hash-coded search for the process entry rather than depending on the offset coded in the process_id.

Fixed on 6180 system:

<u>problem name</u>	<u>nature of problem</u>	<u>fix</u>
ii and fim assume locked sb	Both the interrupt interceptor and fault interceptor module assume that the sb register is locked. Since it is not, the user can load sb with special values and cause overwriting of any ring 0 data segment at the next fault or interrupt.	short-term: add check to ii and fim for ring 0 before believing sb. long-term: eliminate need to depend on sb by simplifying interrupt and fault handling.
New ring stack bug	Gatekeeper, upon creating a new ring, is willing to use a previously existing stack segment, which may have incorrect ring brackets, thus exposing an inner ring stack to an outer ring.	Gatekeeper should check error code returned by makeseg to see if a new stack segment was created. If not, should signal error.
Syserr masking too long	When an error inside the system occurs, the syserr routine masks the CPU against interrupts while printing a message on the operator's console. If too many interrupts come in, GIOC or IOM status queues will overflow, crashing system. User can trigger an apparent system error by producing an op-not-complete fault; he can then generate enough interrupts to crash the system by sending a stream of characters with incorrect parity.	Syserr printing routine should be revised to permit interrupts to be handled normally during message printing. 645F CPU will eliminate user ability to generate op-not-complete faults.
GIM data base bug	The GIM creates a data base by a call to "makeseg" rather than append branch. As a result, it will use any segment which is already around and which has the right name. User can then overwrite the GIM data.	Fix GIM to call append branch rather than makeseg.

<u>problem name</u>	<u>nature of problem</u>	<u>fix</u>
unvalidated gates	No argument validation is specified for the following gates to ring zero: absentee_test_ (all entries) hphcs_ (all entries) phcs_ (all entries) phnxhcs_ (all entries) admin_gate_\$guaranteed_eligibility_off admin_gate_\$guaranteed_eligibility_on	Wait for 6180
Incorrectly validated arguments	In the following entries, some argument is validated with more leniency than appropriate, permitting the user, typically, to cause the supervisor to write into an area in which the user has no access. hcs_\$get_seg_count last argument unvalidated. hcs_\$get_entry_name argument validated for wrong type. hcs_\$get_dbrs argument validated for wrong usage. hcs_\$assign_channel 1st argument validated for wrong usage. hcs_\$check_device 2nd argument validated for wrong usage. hcs_\$get_search_rule argument validated for wrong usage. hcs_\$get_count_linkage 2nd argument validated for wrong usage. hcs_\$ipc_init argument validated for wrong usage. hcs_\$list_dir 2nd argument validated for wrong usage. hcs_\$make_ptr 1st argument validated for wrong usage. hcs_\$list_dir_acl 3rd argument validated for wrong usage. hcs_\$set_dtd 3rd argument validated for wrong usage. hcs_\$status entire argument spec is wrong. imp_dim_gate_\$imp_read_order 3rd argument validated for wrong usage. imp_dim_gate_\$imp_write_order 3rd argument validated for wrong usage. netp_\$ncp_priv_status 3rd argument validated for wrong usage. netp_\$ncp_priv_order 3rd argument validated for wrong usage. net_\$ncp_status 3rd argument validated for wrong usage. net_\$ncp_order 3rd argument validated for wrong usage. hcs_\$acl_list 5th argument validated for wrong usage.	Wait for 6180

<u>problem name</u>	<u>nature of problem</u>	<u>fix</u>
unvalidatable arguments	In the following entries, some entry <u>cannot</u> be checked by the automatic validator, since the correct method of validation depends on the value of some other argument.	Wait for 6180 (Calling sequence should eventually be changed, also.)
	hcs_\$acl_list	3rd argument used as both input and output.
	hcs_\$ex_acl_list	3rd argument used as both input and output.
	hcs_\$ex_acl_delete	3rd argument meaning depends on 4th argument.
	hcs_\$initiate_seg_count	6th argument meaning depends on another argument.
	hcs_\$list_dir_acl	4,5th arguments meaning depend on the value of 3rd argument.
	hcs_\$replace_sall	3rd argument unvalidatable.
	hcs_\$replace_dall	
EPL argument validation trap	The argument validator does not completely check out certain EPL specifiers. When combined with an incorrectly specified gate, can allow misuse of that gate.	Wait for 6180
CPU hogging bug #3	Entry point hcs_\$list_connect will, if called with proper argument, set the "interaction switch" on, giving user credit for an interaction the next time he calls block. By calling every few seconds, one can stay forever in the highest scheduling queue, and completely deny service to lower queue users.	Review PIM design (GIM replacement) to make sure it does not have same hole.