



SMASER

Questions and Answers

TRIPsystem
Product Documentation



End User License Agreement

All rights to this software, its documentation and logotypes of the TRIP product family and software (altogether "Software") supplied by Smaser AG (Smaser) are exclusively owned by Smaser.

The transfer of this Software, solutions or parts thereof requires the prior written agreement of Smaser. Furthermore, the customer has the right to use licensed Software and / or process solutions supplied by Smaser to the extent specified in his contract with Smaser.

The free-to-use non-commercial version doesn't require a prior written agreement with Smaser but such customers, organizations and/or third parties agree by using the software and / or solution of Smaser to be strongly obliged to keep all rights to this software, documentation and logotypes of the TRIP product family absolutely unfringed and protected.



Table of Contents

PREFACE.....	4
THE QUESTIONS AND ANSWERS	4
Why does the deletion of a part record change the part record number for each of the part records following the one deleted?	4
Is there a date format available that TRIP always accepts?	4
Is it possible to obtain the names of databases belonging to a cluster?	4
Can you specify an alternate format that can be used with <i>Show FOCUS</i> ?	4
Are return codes provided from a called macro to the one, which called it?	4
Is the <i><sortfields></i> function in an output format ignored when the database associated with it is a member of a cluster?	5
When does TRIP remove the hyphen for indexing? How is it used for searching?	5
Why is the order of the output not intuitive with clustered databases?	5
What commands can be executed in TRIP's <i>TTY</i> mode?	5
Can you apply a 'first word' wildcard at the end of a search pattern?	5
Is it possible to get the list of terms from a clustered database in a Data Entry form using <i><PF1> <D></i> ?	6
Is there a command to limit the number of records printed?	6
Is it possible to specify <i>DEfine Print MAXimum</i> for <i>Print Local</i> ?	6
Am I allowed to change the system time while TRIP is started?	6
Is there a limit for a search executed either in a procedure or in <i>CCL</i> ?	6
Is the return code of the ASE routine <i>CheckMode</i> unique?	6
Are changes to cluster(s) or databases valid for other TRIP sessions, which have been opened?	6
Should I restart TRIP to use a modified cluster definition?	6
Can different users get mail when an indexing error occurs?	7
What happens if the available system memory resources are exhausted?	7
Being logged in as the TRIP-user <i>SYSTEM</i> , can I get a list of all existing users in TRIP?	7
Can I index a TRIP 2.4 database with TRIP 3.x?	7
How can a search result be incorporated in the <i>Print FRequency</i> command?	7
How does one move to the end of a long list of record parts during data entry?	7
Is there a limit on highest record number and the number of records in a TRIP database?	7
Can one use Search Forms from ASEs?	8
What values must be defined if I want to have more than 130 users in parallel on a VMS system?	8
Are you aware of any problems when the physical size of the database gets too big?	8
Is there a way to find out which user damaged the <i>BAF</i> file?	8
How can I combine a variable and filename in a macro to define a complete path for a file operation?	8
I am writing a TRIPapi program on UNIX and I wish to use TRIP's interrupt handlers to generate the equivalent of the <i>Searching...</i> and <i>Still searching...</i> comforters. Everything works fine as long as I have a terminal attached to my session, but as soon as I run as a detached server, it stops working. How can I make it work again?	8



Preface

This document contains answers to some questions, relating to TRIPsystem, that are not already covered in other TRIPsystem documentation.

The Questions and Answers

Why does the deletion of a part record change the part record number for each of the part records following the one deleted?

There are two series of numbers used in TRIP for the part records of a given record. One series – used internally by TRIP, not available to the user - consists of the nominal numbers of the part records. The other series – the one available to the users, e.g. shown in output formats – consists of the ordinal numbers of the parts, indicating how the parts are stored within the record: the first part is presented as part record 1, the second as part record 2 etc.

It follows that, for the numbers seen by the users, the number of a given part can never be greater than the total number of parts – and so, if a part is deleted then the remaining parts (as seen by the users) will be renumbered (exactly how depends on where in the sequence of parts the deleted one resided).

If an application requires that a given part record keeps the same number permanently then no part must ever be deleted and new parts must not be inserted between existing parts. N.B. A part is not deleted just because all its contents are deleted. An alternative method of giving the part records in a record with unique and permanent identities is, of course, to design the database with a part record name field.

(In earlier versions of TRIP, some aspects of output formatting did, in certain contexts, incorrectly present a part record number from the nominal number series to the users. These instances have been corrected in V3.2-1, cf. Chapter 2).

Is there a date format available that TRIP always accepts?

With TRIP 3.0 and later, the date form YYMMDD is no longer standard or a universal default format. The users must alter their own specified formats (user profiles) to produce dates using this pattern. The date format YYYYMMDD is, however, universally accepted as long as all eight digits are specified when searching. When performing truncated date searching, such as Find 1993, the user's profile is used to determine whether this means YYYY or YYMM.

There is a default CCL reference language file shipped with TRIP, e.g. TDBS\$SYS:ENGLISH.CCL. This default will be in effect until a user modifies his/her own user profile. Generally, you should check your date format definition in your profile. TRIP 3.1 will now save the date format in the history as it was entered.

Is it possible to obtain the names of databases belonging to a cluster?

In CCL, the STatus command on a cluster name will list, as part of the report, the databases which are members of that cluster. Programmatically, the TRIPapi function CurrentBase will write the database names, separated by commas, into the shin element of the SHELL/DAO interface record.

Can you specify an alternate format that can be used with *Show FOCUS*?

Yes. Use the CCL command: Show FOCUS Format=formatname

Are return codes provided from a called macro to the one, which called it?

No. No major enhancements have been made to the MACRO language for TRIP.



Is the <sortfields> function in an output format ignored when the database associated with it is a member of a cluster?

The identifier is not ignored, but it is not handled as might be expected. The <sortfields> field parameters are translated to field numbers, which are then accumulated together so that the field numbers used in all of the <sortfields> statements, used in the combined output formats, will be applied to all databases in the cluster. If the database designs are not identical, the results will be 'interesting' at best. We recommend that <sortfields> not be used in output formats used in commonly clustered databases unless those databases have similar structure.

When does TRIP remove the hyphen for indexing? How is it used for searching?

As long as hyphen is not specified in the database design as being a special searchable character, hyphens are removed (for indexing purposes, not from the actual record in the BAF file of the database) when they occur as the last character of a line and the first character of the next line is in lower case; and they will then also be handled as space characters in a search. You will therefore find that, for example, Find mouse-hole will identify all occurrences in the database of mouse hole and mouse-hole (and upper and lower case variants thereof) except those cases of mouse-hole where the hyphen is the last character of a line in the database (and the h character - being the first one on the next line - is not in upper case).

Why is the order of the output not intuitive with clustered databases?

If two databases are opened in the following order:

- BAsE database1
- BAsE database2
- BAsE temp = database2,database1

why does the output from the cluster still come out with database1 first?

During the current session and including any carryover from a saved .SIF file, the order in which output is generated for clustered databases, is determined by the original order in which the individual databases within the cluster were opened at any time. When opened, a database receives an internal index, which is kept during the whole session - or until the CClose command has been used on it, and other restarted sessions. In the example above, since database1 was originally opened first its output will appear first, unless subjected to a SORT MERGE.

If you wish to ensure a certain order of database output, use a startup procedure to open the desired databases sequentially before the user has the opportunity to change that order - and, please note that, if the user applies the CClose command for all open databases then the order he later restarts the databases will decide the order in which the output from them is going to be given.

What commands can be executed in TRIP's TTY mode?

Let's give the answer the other way around: You cannot execute any CCL commands that would require a full screen operating. i.e. EDit, EDit Copy, EDit INSert, or selecting from a Show or Display list.

Can you apply a 'first word' wildcard at the end of a search pattern?

No. The following search commands are not valid in TRIP:

- Find chapter=& alice &
- Find chapter=& al\$ s evidence &

Both commands result in an Illegal order part message. However, a search for & ALICE, or & AL\$ s evidence is correct.



Is it possible to get the list of terms from a clustered database in a Data Entry form using <PF1> <D>?

Yes and no. TRIP 3.0 and later permits the use, during field definition, of multiple Database.Field references when defining a restriction or controlled term list (during database creation). This will result in terms from the corresponding databases and fields to be displayed together when <PF1> <D> is pressed during data entry. However, it is not possible to apply the technique against a dynamically defined database cluster.

Is there a command to limit the number of records printed?

Yes. Use the CCL command:

- DEfine Print MAXimum=n
where n is the maximum number of records allowed to print.

Is it possible to specify *DEfine Print MAXimum* for *Print Local*?

No, Print Local is not affected by the DEfine Print MAXimum command.

Am I allowed to change the system time while TRIP is started?

No, especially if you run TRIP batch jobs such as Index or Global Update. You must not change the time whenever a user is working with the system. TRIP keeps many time stamps for each database, therefore even when nobody is actually working with the system, rolling back the system time in most cases will result in TRIP simply disallowing any further logins to the system. It is possible to set the system time/date ahead.

Is there a limit for a search executed either in a procedure or in CCL?

Yes. You cannot perform a search command that has more than 400 characters. You are allowed to save long find commands while editing a procedure (e.g. to split it up later), but you will get an error message when trying to run the procedure in CCL (CCL command line 'n' in 'procedurename' too long).

A new feature of TRIP 3.0 is the ability to pick up as many terms from a Display list as desired. A search incorporating these terms will be executed properly: It splits the search into several OR-ed search commands. Therefore we recommend proceeding in such a way with all long search statements.

If you are creating a new application (or rewriting an old one), you can take advantage of the new functionality of TRIP that the limit of 400 characters in a CCL order is now removed. To achieve this, TRIP must be started with the new API function TdbStartTrip() instead of the old one, TdbBeginTdb(), and you must use the new API functions to handle the CCL order (see above).

NB: TRIPclassic will continue to call TdbBeginTdb() when starting up so it cannot benefit from this new feature.

Is the return code of the ASE routine *CheckMode* unique?

Not in all cases. The routine returns the same return code for both EDit Copy and Data Entry - Add. In all other cases, the return code is unique.

Are changes to cluster(s) or databases valid for other TRIP sessions, which have been opened?

No. The TRIP user who wants to see or work with the last cluster or database update must reopen the database. The user must restart TRIP if the design has been changed.

Should I restart TRIP to use a modified cluster definition?

Yes. A cluster (re)definition is handled in exactly the same manner as a database (re)definition. To ensure that the change takes effect in your own TRIP session, you must exit



the session and restart TRIP. We strongly advise you do this, otherwise there will be inconsistent behaviour within either the cluster or the database.

Can different users get mail when an indexing error occurs?

Yes. You can specify the logical TDBS\$ERRMAILST (UNIX: TDBS_ERRMAILST) specifying a list of username's to receive mail whenever there has been an index or global update error. For example:

- VMS : DEFINE TDBS\$ERRMAILST USR1,USR2,USR3
- UNIX (ksh) : export TDBS_ERRMAILST="USR1 USR2 USR3"

The mail message is generated and sent to the users specified for either an unsuccessful Index or a failed Global Update command.

NB: On UNIX, separate the names with spaces; on VMS separate them with commas.

What happens if the available system memory resources are exhausted?

This question arose when users performed CCL commands involving TRIP to request a large amount of working memory. E.g. a Next 50000 may result in a crash when the system is out of memory. Due to performance aspects TRIP does not check when the available working-set gets exhausted. Therefore, exceeding the amount of available memory results in a system crash.

Being logged in as the TRIP-user **SYSTEM**, can I get a list of all existing users in TRIP?

Yes, but the CONTROL-database must have been indexed after a user manager added the users.

Can I index a TRIP 2.4 database with TRIP 3.x?

No. TRIP V3.1-2 or later no longer supports old style indices—you must upgrade your index files before attempting to use V3.1-2 or later (see the TRIP System Manager's Guide for more detail).

How can a search result be incorporated in the **Print F**requency command?

If a Print FFrequency command is to include a search number then it must be entered in the command string after Print FFrequency and before fieldname. If the search number is entered after fieldname, then TRIP will identify it as an illegal order; however if entered after Print, but before FFrequency, TRIP will interpret it as a Step value. The command will be accepted but subsequently fail for constituting an illegal Step / Range specification. The best way, therefore, is to specify Print FFrequency commands in the following order:

Print FFrequency S=n fieldname

How does one move to the end of a long list of record parts during data entry?

In data entry <PF1> <N> and <PF1> <P> are used to move between record parts. To quickly access the last record part, first press <PF1><PF12> (VMS) or <PF1><Ctrl> (UNIX), then press <PF1> <P> (if the record does not include a part name field) or <PF3> (if the record includes a part name field).

Is there a limit on highest record number and the number of records in a TRIP database?

The actual limit, for both highest record number and the number of records in TRIP, is 2^{32} (2 to the 32nd power—approximately 4 billion). Problems could occur with the maximum number of records if, for instance, you have 2^{30} records in your database, the number of occurrences for a given term could exceed the 2^{32} limit by having more than 4 occurrences per record (e.g. the word 'the').



For the maximum record number, TRIP will create a record with a higher record number than 2^{32} , but the TRIPapi will not be able to deliver it to an application program (such as TRIPclassic). This situation should never occur unless the number of records also reaches such a high value.

Can one use Search Forms from ASEs?

No. You can neither call a Search form nor any other TRIPclassic form from within an ASE.

What values must be defined if I want to have more than 130 users in parallel on a VMS system?

The rule should be:

- $ASTLM = 1$ for every concurrent user + 10%
- $FILLM = 2$ for every concurrent user + 10%

E.g with 130 users:

- $ASTLM = 1 (130) + 10\% (13) = 143$ (round up to 150)
- $FILLM = 2 (260) + 10\% (26) = 286$ (round up to 300)

Are you aware of any problems when the physical size of the database gets too big?

As a database gets larger, it gets slower. Literal term searches will continue to be the same speed regardless of the size of the database, so long as the underlying hardware / operating system can support the physical size of the BIF file. Computed searches however will get slower. A good size to aim for is 2 GB. If you intend to build a large database, you should consider clustering several physical databases rather than simply building one large file.

Is there a way to find out which user damaged the BAF file?

With TRIP 3.1, when the message The BAF file of 'databasename' needs reorganisation appears, press <PF2> and you will see the TRIP user who modified the BAF file last. There is no information stored regarding modification to modification as to what exactly was modified.

How can I combine a variable and filename in a macro to define a complete path for a file operation?

Create a procedure that would prompt for a directory, then spawn a DCL command that would perform a logical name definition in the user's job table. This logical could then be used in future filename definitions within the macro.

I am writing a TRIPapi program on UNIX and I wish to use TRIP's interrupt handlers to generate the equivalent of the *Searching...* and *Still searching...* comforters. Everything works fine as long as I have a terminal attached to my session, but as soon as I run as a detached server, it stops working. How can I make it work again?

TRIP explicitly switches off its comforters if it detects that there is no terminal, i.e. it thinks that the user is operating in batch mode and therefore does not need "comforting" during long searches. The way in which TRIP determines its running mode is by using the `isatty()` function and programmers who require comforters in detached processes can simply override this function with their own definition. In this case, you will need to always return a success value (1) from your function "`isatty(int)`", for example:



```
int isatty(int ttyDescriptor)
{
    return 1;
}
```

Then, the link instruction for your program can explicitly include this function before the default C library, libc.