

Running the Universal File Interface (UFI)

NetCDF Demo on VMware

Introduction

Many “technologists” (e.g., scientists, engineers, and other researchers) avoid using databases; a principal reason given is that they often have very large data files which are too cumbersome/difficult/slow/costly to load into a database. These files come in a variety of formats: [HDF5](#), [NetCDF](#), [NITFS](#), [FITS](#), etc. At BCS we have developed a solution to this problem: the Universal File Interface (UFI), a database extension based on the IBM Informix Virtual Table Interface (VTI).

VTI is a technology that supports making external datasets appear as tables to SQL queries and statements¹. UFI is a BCS database extension for delivering the contents of external data files as though they were rows in a database table. UFI makes a file look like a set of database tables, so “UFI-managed tables” are actually *virtual* database tables. Consequently, at no point do users of UFI need to load their files into a database!

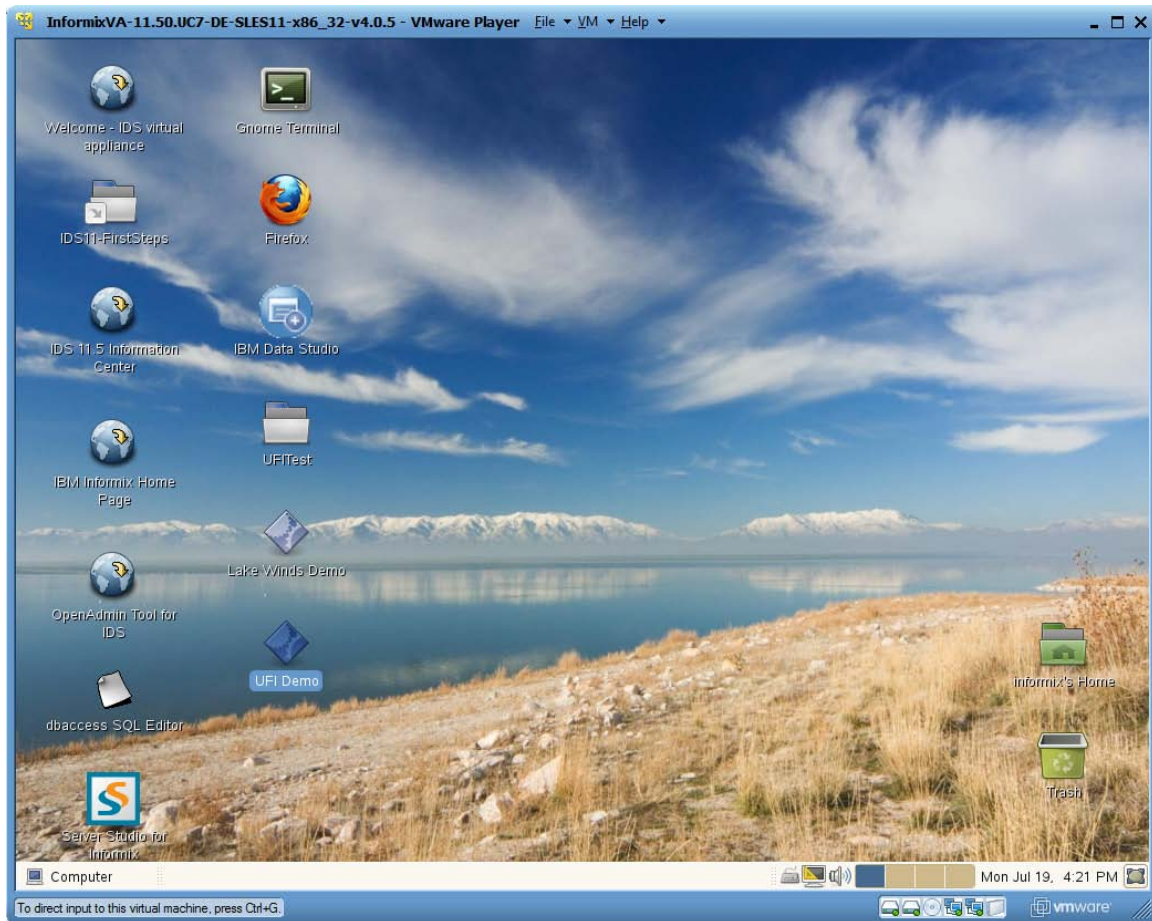
A trial copy of UFI can be requested from BCSInfo@barrodale.com. If you install UFI then you can run a demo program that we have supplied to illustrate some of the key features of UFI. If you would prefer to just read about this demo, then the remainder of this document shows you what you would see if you ran the demo.

Running the Demo

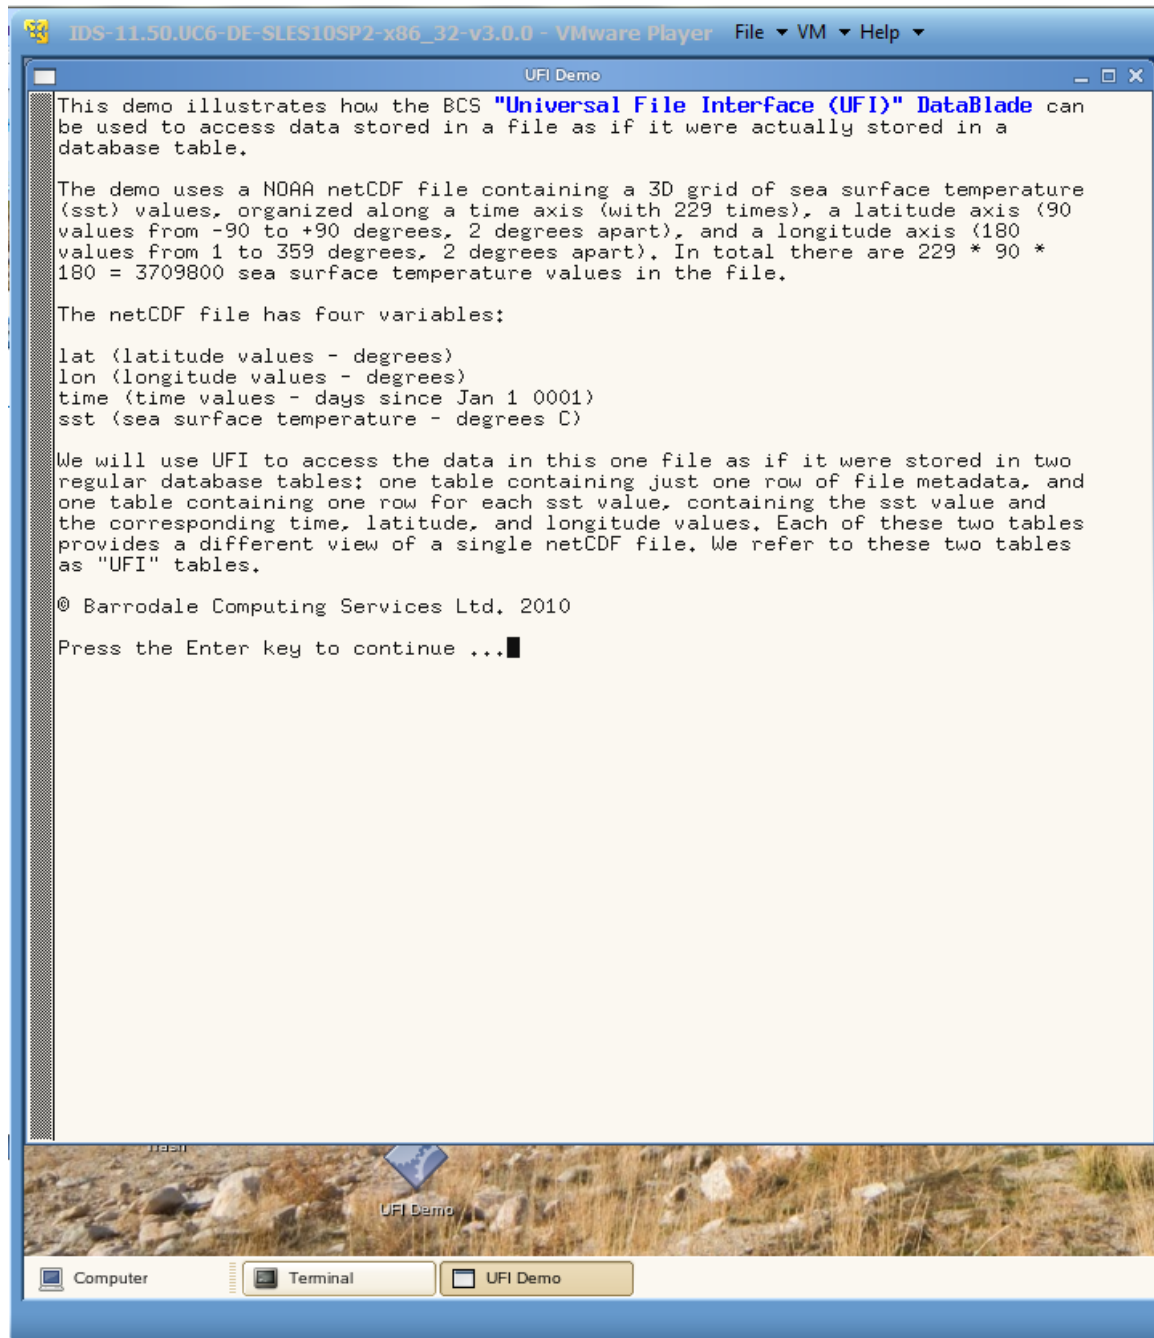
Start the IDS Developer Edition virtual appliance as described in “[The Informix / Universal File Interface \(UFI\) VMware Appliance](#)” (which can be downloaded from http://www.barrodale.com/bcs/downloads/ufi_ids_appliance/UFI_InstallVMware.pdf).

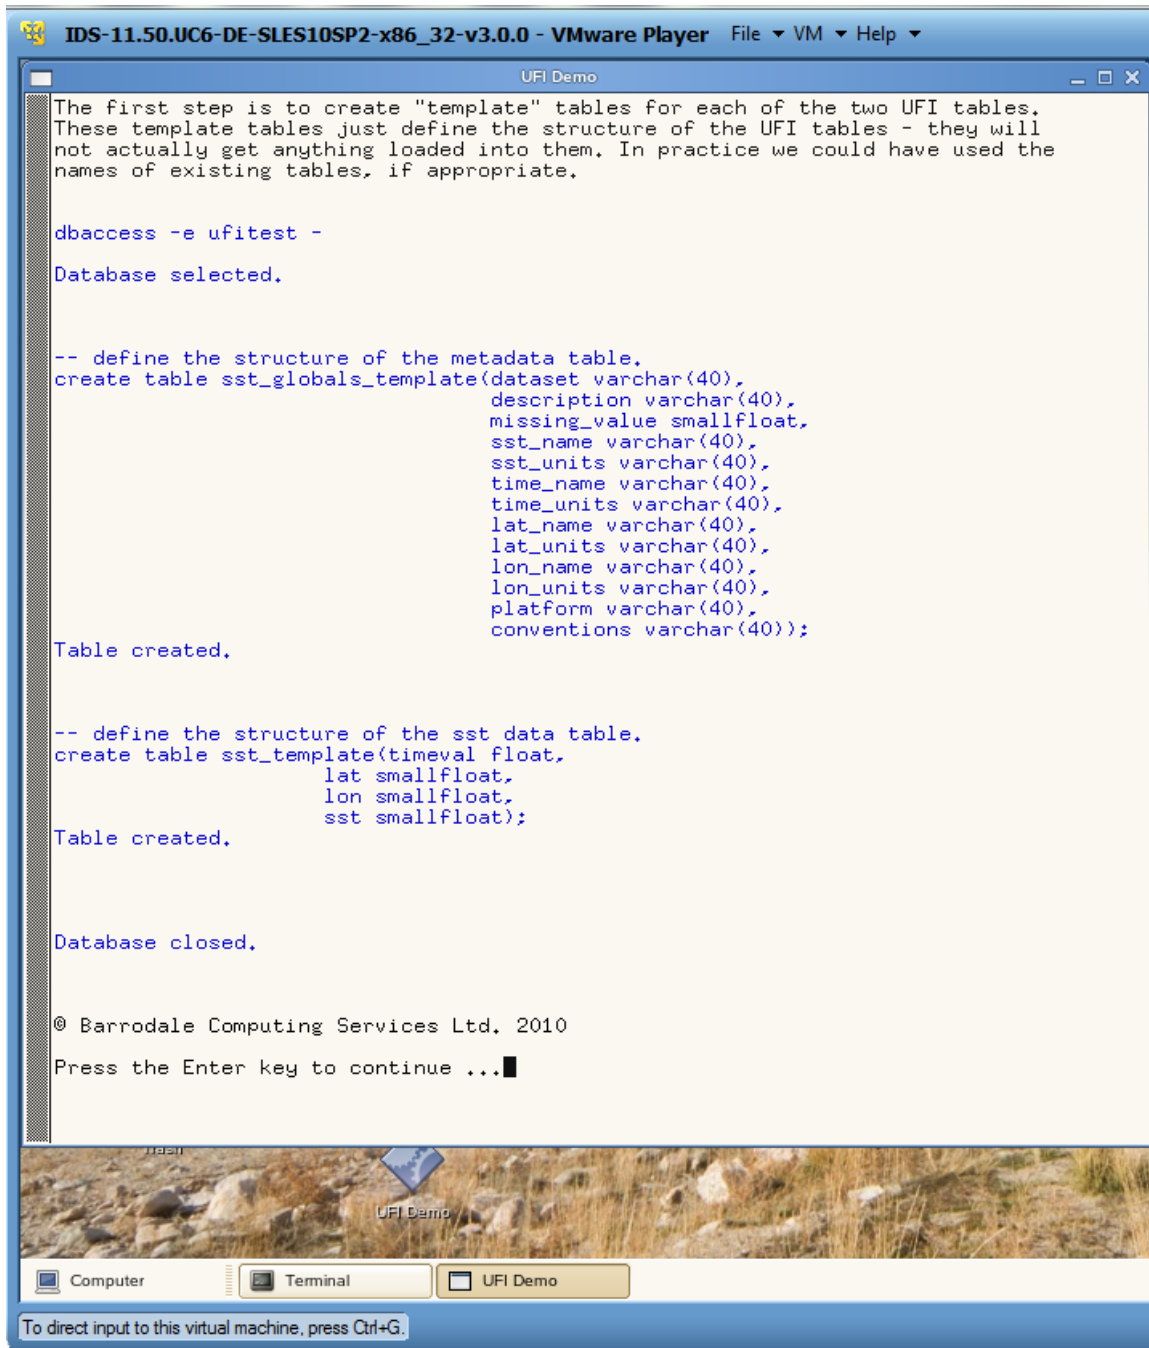
¹ It addresses the same needs as the SQL/MED, or Management of External Data, extension to the SQL standard as defined by ISO/IEC 9075-9:2003.

Log in as user “informix” (password “informix”) and locate the “EFI Demo” icon (highlighted in the following picture):



Double-click on the “EFI Demo” icon. You will then see the following sequence of pages.





```
IDS-11.50.UC6-DE-SLES10SP2-x86_32-v3.0.0 - VMware Player  File ▾ VM ▾ Help ▾
UFI Demo
The first step is to create "template" tables for each of the two UFI tables.
These template tables just define the structure of the UFI tables - they will
not actually get anything loaded into them. In practice we could have used the
names of existing tables, if appropriate.

dbaccess -e ufitest -
Database selected.

-- define the structure of the metadata table.
create table sst_globals_template(dataset varchar(40),
                                description varchar(40),
                                missing_value smallfloat,
                                sst_name varchar(40),
                                sst_units varchar(40),
                                time_name varchar(40),
                                time_units varchar(40),
                                lat_name varchar(40),
                                lat_units varchar(40),
                                lon_name varchar(40),
                                lon_units varchar(40),
                                platform varchar(40),
                                conventions varchar(40));

Table created.

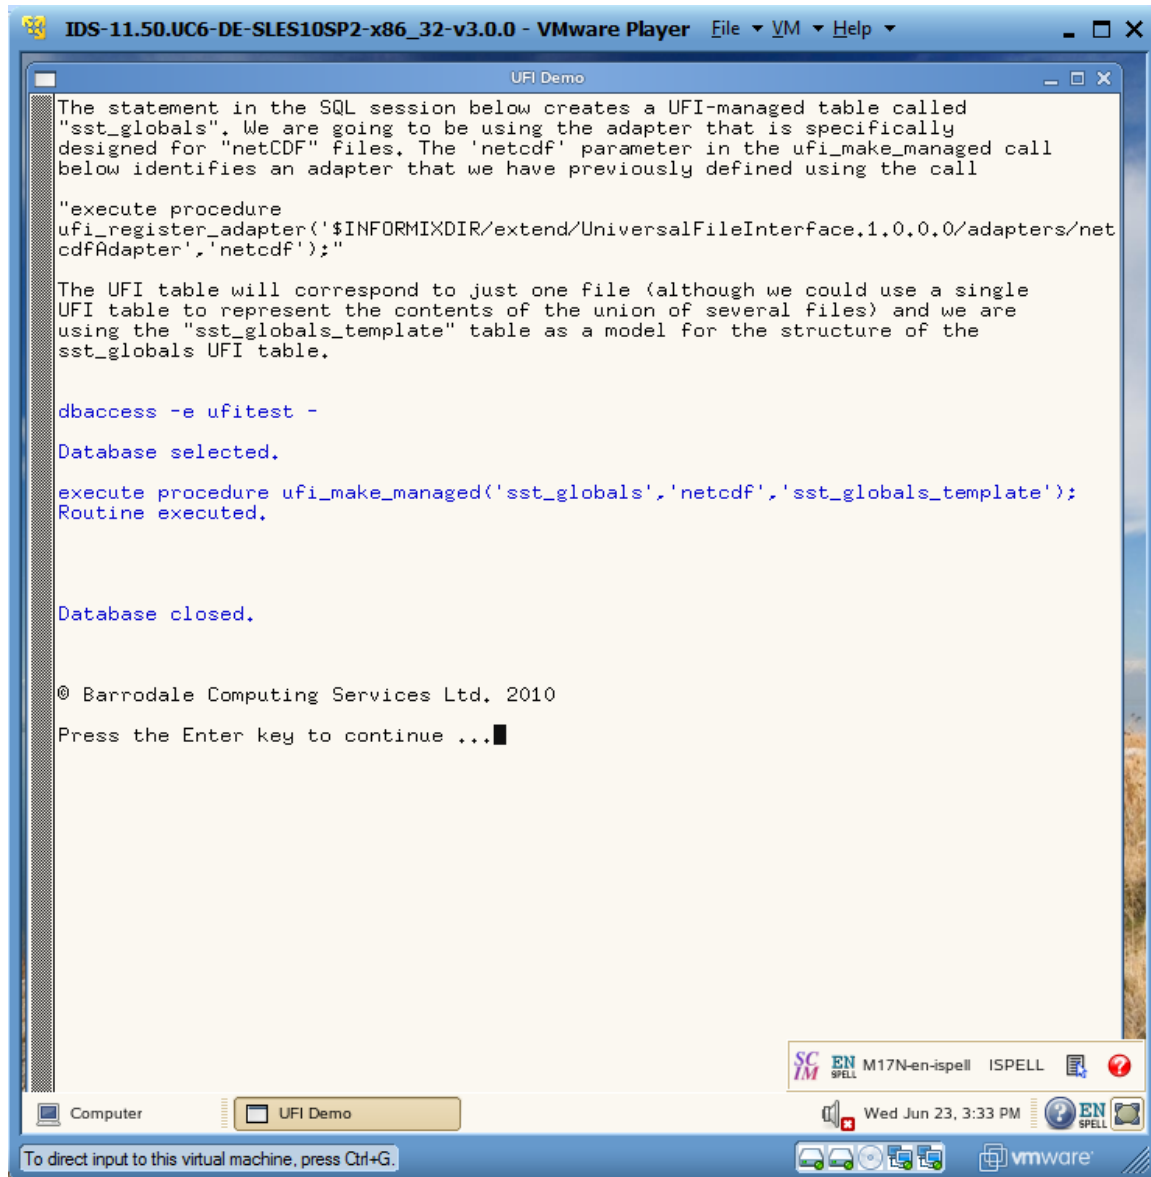
-- define the structure of the sst data table.
create table sst_template(timeval float,
                          lat smallfloat,
                          lon smallfloat,
                          sst smallfloat);

Table created.

Database closed.

© Barrodale Computing Services Ltd. 2010
Press the Enter key to continue ...█

Computer  Terminal  UFI Demo
To direct input to this virtual machine, press Ctrl+G.
```



IDS-11.50.UC6-DE-SLES10SP2-x86_32-v3.0.0 - VMware Player File VM Help

UFI Demo

The statement in the SQL session below creates a UFI-managed table called "sst_globals". We are going to be using the adapter that is specifically designed for "netCDF" files. The 'netcdf' parameter in the ufi_make_managed call below identifies an adapter that we have previously defined using the call

```
"execute procedure
ufi_register_adapter('$INFORMIXDIR/extend/UniversalFileInterface.1.0.0.0/adapters/netcdfAdapter','netcdf');"
```

The UFI table will correspond to just one file (although we could use a single UFI table to represent the contents of the union of several files) and we are using the "sst_globals_template" table as a model for the structure of the sst_globals UFI table.

```
dbaccess -e ufitest -
Database selected.
execute procedure ufi_make_managed('sst_globals','netcdf','sst_globals_template');
Routine executed.

Database closed.
```

© Barrodale Computing Services Ltd. 2010
Press the Enter key to continue ...

Computer UFI Demo

SC EN M17N-en-ispell ISPELL
Wed Jun 23, 3:33 PM EN SPELL
vmware

To direct input to this virtual machine, press Ctrl+G.



The statements in the SQL session below each add a single column to the "sst_globals" UFI table that we just created. In total, there will be 13 columns in this UFI table.

The first parameter in each call is the UFI table name. The second parameter in each call is the UFI column name. The third argument in each call contains the name of the netCDF attribute or variable that we are mapping to the UFI table column.

In the first five calls we are mapping attributes of the sst netCDF variable. In the next six calls we are mapping attributes of the time, lat, and lon netCDF variables. In the last two calls we are mapping global attributes.

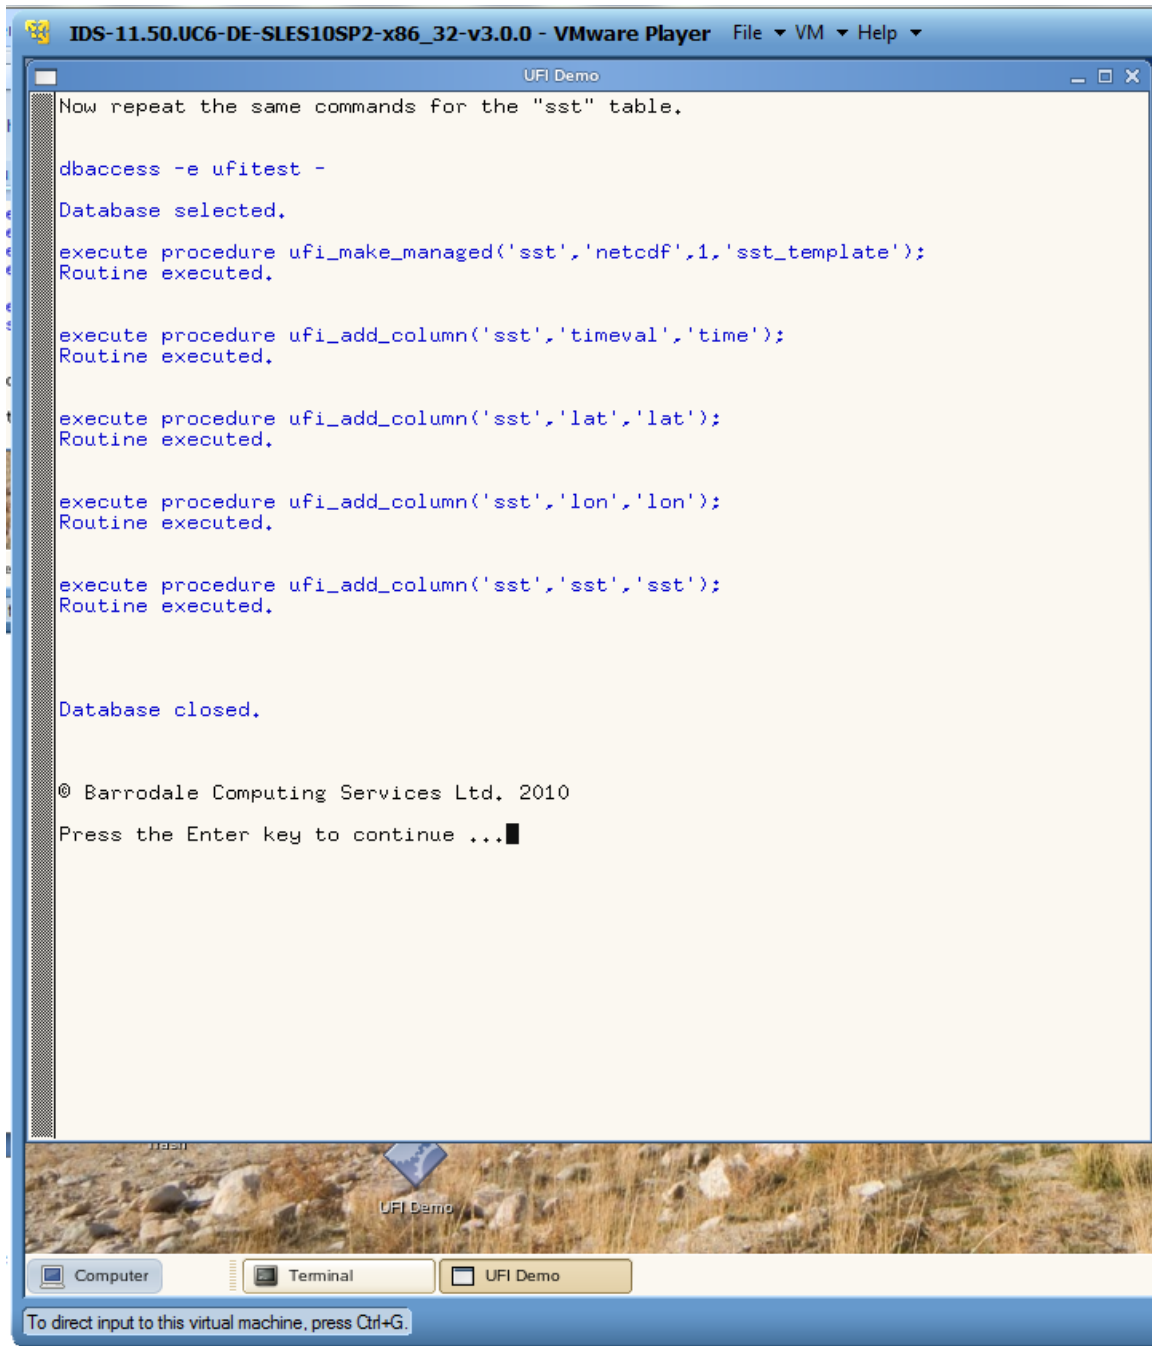
```
dbaccess -e ufitest -
Database selected.
execute procedure ufi_add_column('sst_globals','dataset','sst.dataset');
Routine executed.
execute procedure ufi_add_column('sst_globals','description','sst.var_desc');
Routine executed.
execute procedure ufi_add_column('sst_globals','missing_value','sst.missing_value');
Routine executed.
execute procedure ufi_add_column('sst_globals','sst_name','sst.long_name');
Routine executed.
execute procedure ufi_add_column('sst_globals','sst_units','sst.units');
Routine executed.
execute procedure ufi_add_column('sst_globals','time_name','time.long_name');
Routine executed.
execute procedure ufi_add_column('sst_globals','time_units','time.units');
Routine executed.
execute procedure ufi_add_column('sst_globals','lat_name','lat.long_name');
Routine executed.
execute procedure ufi_add_column('sst_globals','lat_units','lat.units');
Routine executed.
execute procedure ufi_add_column('sst_globals','lon_name','lon.long_name');
Routine executed.
execute procedure ufi_add_column('sst_globals','lon_units','lon.units');
Routine executed.
execute procedure ufi_add_column('sst_globals','platform','nc_globals.platform');
Routine executed.
execute procedure ufi_add_column('sst_globals','conventions','nc_globals.Conventions');
Routine executed.
Database closed.
```

© Barrodale Computing Services Ltd, 2010

Press the Enter key to continue ...

Computer Terminal UFI Demo

To direct input to this virtual machine, press Ctrl+G.



```
IDS-11.50.UC6-DE-SLES10SP2-x86_32-v3.0.0 - VMware Player  File ▾ VM ▾ Help ▾
UFI Demo
Now repeat the same commands for the "sst" table.

dbaccess -e ufitest -
Database selected.
execute procedure ufi_make_managed('sst','netcdf',1,'sst_template');
Routine executed.

execute procedure ufi_add_column('sst','timeval','time');
Routine executed.

execute procedure ufi_add_column('sst','lat','lat');
Routine executed.

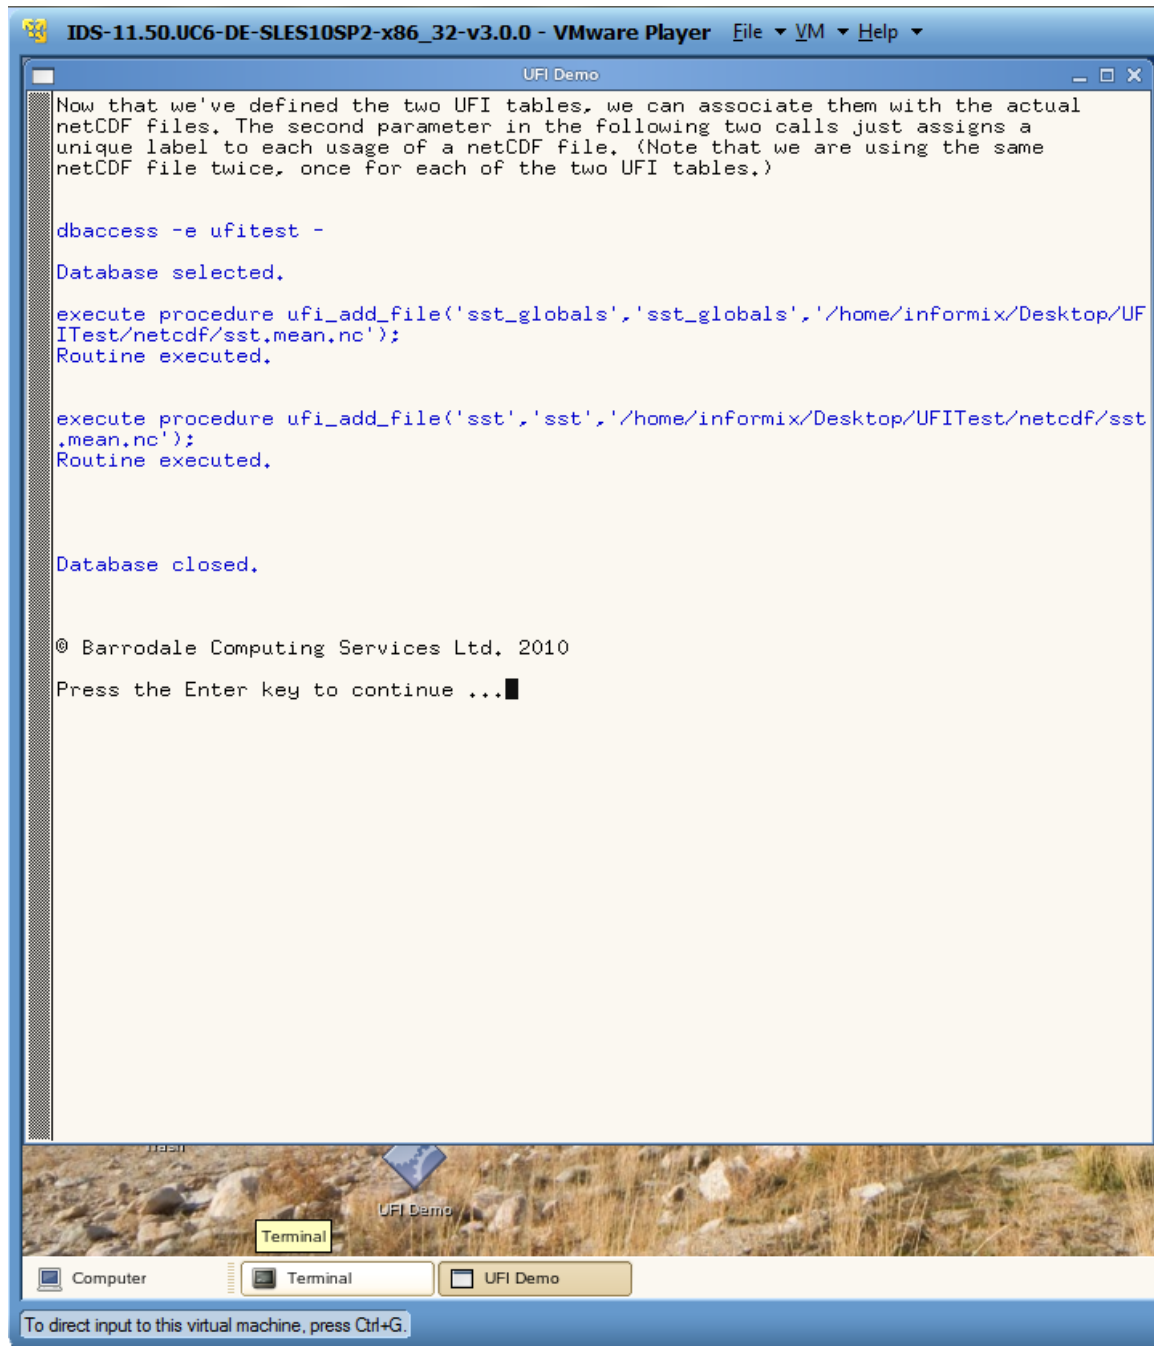
execute procedure ufi_add_column('sst','lon','lon');
Routine executed.

execute procedure ufi_add_column('sst','sst','sst');
Routine executed.

Database closed.

© Barrodale Computing Services Ltd. 2010
Press the Enter key to continue ...■

Computer  Terminal  UFI Demo
To direct input to this virtual machine, press Ctrl+G.
```



The screenshot shows a VMware Player window titled "IDS-11.50.UC6-DE-SLES10SP2-x86_32-v3.0.0 - VMware Player". Inside the player, a terminal window titled "EFI Demo" is open, displaying the following text:

```
Now that we've defined the two EFI tables, we can associate them with the actual netCDF files. The second parameter in the following two calls just assigns a unique label to each usage of a netCDF file. (Note that we are using the same netCDF file twice, once for each of the two EFI tables.)

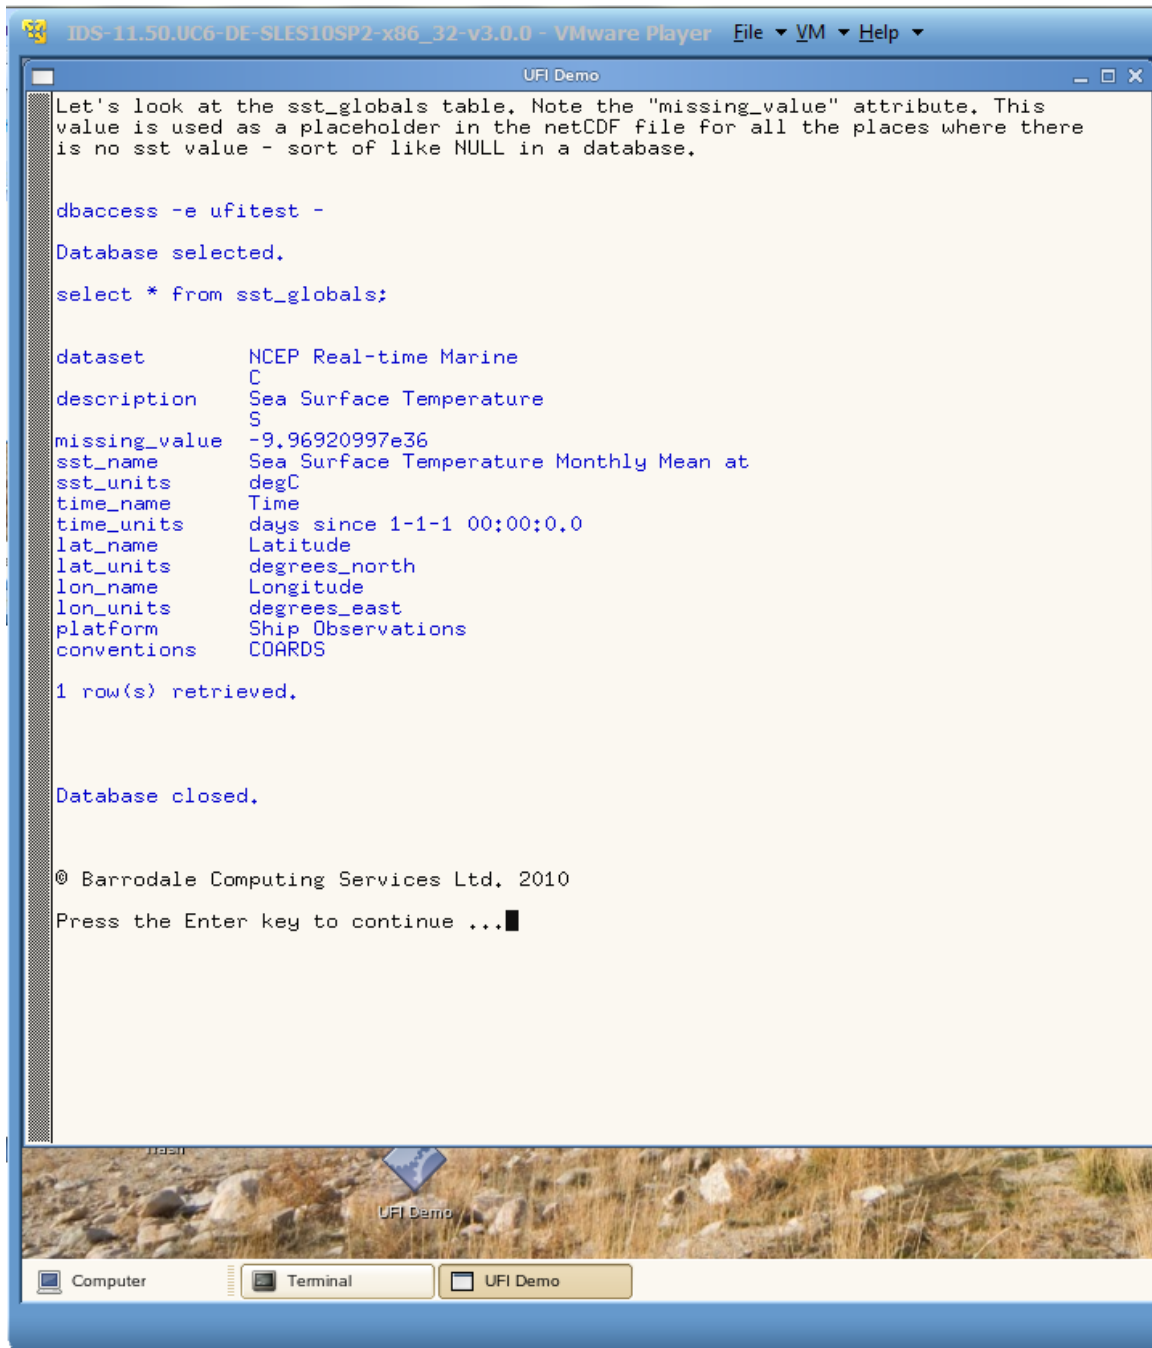
dbaccess -e ufitest -
Database selected.
execute procedure ufi_add_file('sst_globals','sst_globals','/home/informix/Desktop/UFITest/netcdf/sst.mean.nc');
Routine executed.

execute procedure ufi_add_file('sst','sst','/home/informix/Desktop/UFITest/netcdf/sst.mean.nc');
Routine executed.

Database closed.

© Barrodale Computing Services Ltd. 2010
Press the Enter key to continue ...
```

At the bottom of the VMware Player window, there is a taskbar with three buttons: "Computer", "Terminal", and "EFI Demo". Below the taskbar, a message reads: "To direct input to this virtual machine, press Ctrl+G."



IDS-11.50.UC6-DE-SLES10SP2-x86_32-v3.0.0 - VMware Player File VM Help

UFI Demo

Let's look at the sst_globals table. Note the "missing_value" attribute. This value is used as a placeholder in the netCDF file for all the places where there is no sst value - sort of like NULL in a database.

```
dbaccess -e ufitest -  
Database selected.  
select * from sst_globals;  
  
dataset          NCEP Real-time Marine  
C  
description      Sea Surface Temperature  
S  
missing_value    -9.96920997e36  
sst_name         Sea Surface Temperature Monthly Mean at  
sst_units        degC  
time_name        Time  
time_units       days since 1-1-1 00:00:0.0  
lat_name         Latitude  
lat_units        degrees_north  
lon_name         Longitude  
lon_units        degrees_east  
platform         Ship Observations  
conventions      COARDS  
  
1 row(s) retrieved.  
  
Database closed.  
  
© Barrodale Computing Services Ltd. 2010  
Press the Enter key to continue ...
```

Computer Terminal UFI Demo



```
IDS-11.50.UC6-DE-SLES10SP2-x86_32-v3.0.0 - VMware Player  File  VM  Help  ▾
UFI Demo
Consider the following queries:

dbaccess -e ufitest -
Database selected.

select missing_value from sst_globals into temp tempmissingvalue;
1 row(s) retrieved into temp table.

select first 10 sst.* from sst,tempmissingvalue where sst <> missing_value;

      timeval          lat          lon          sst
726834.000000 75.000000000000 5.000000000000 2.000000000000
726834.000000 75.000000000000 7.000000000000 2.659999850000
726834.000000 75.000000000000 9.000000000000 3.149999860000
726834.000000 75.000000000000 15.000000000000 6.299999710000
726834.000000 75.000000000000 21.000000000000 4.599999900000
726834.000000 75.000000000000 23.000000000000 2.679999830000
726834.000000 75.000000000000 25.000000000000 3.460000040000
726834.000000 75.000000000000 27.000000000000 3.599999900000
726834.000000 75.000000000000 29.000000000000 3.699999810000
726834.000000 75.000000000000 31.000000000000 3.639999870000

10 row(s) retrieved.

Database closed.

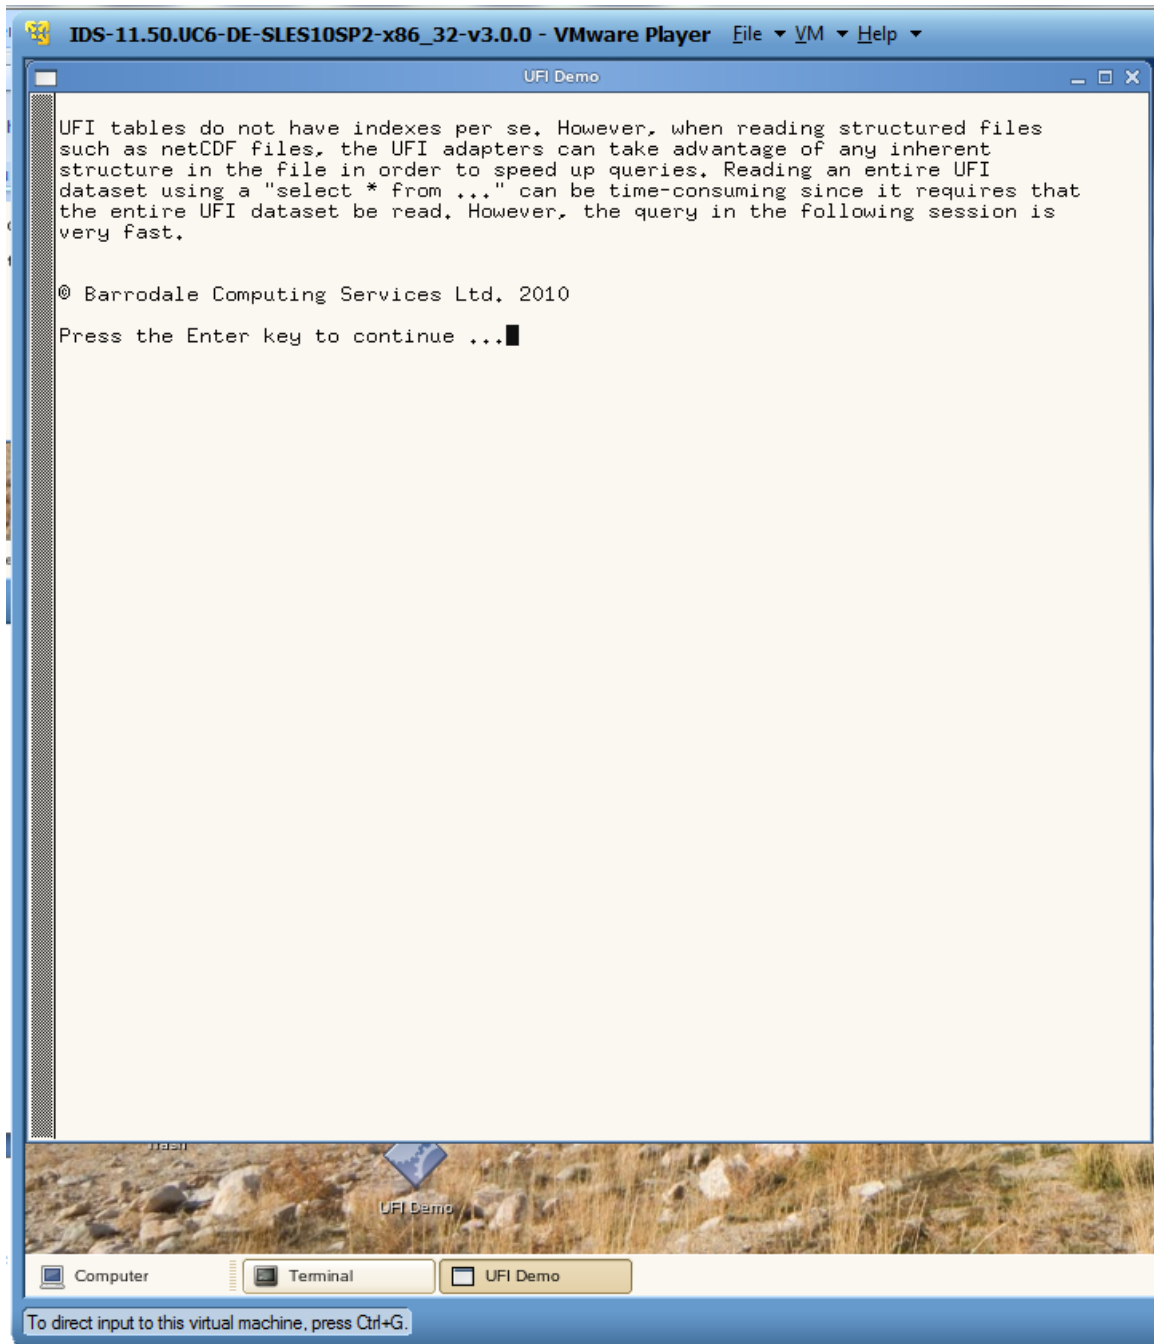
The pair of queries above fetch the first 10 "real" (not "missing") values from
the sst UFI table.

© Barrodale Computing Services Ltd. 2010

Press the Enter key to continue ...
```

UFI Demo

Computer Terminal UFI Demo





The screenshot shows a VMware Player window titled "InformixVA-11.50.UC7-DE-SLES11-x86_32-v4.0.5 - VMware Player". Inside the player is a terminal window titled "UFI Demo" with a menu bar (File, Edit, View, Terminal, Tabs, Help). The terminal displays the following text:

```
dbaccess -e ufitest -  
Database selected.  
select missing_value from sst_globals into temp tempmissingvalue;  
1 row(s) retrieved into temp table.  
  
select current from systables where tabid = 1;  
(expression)  
2010-07-19 16:28:46.000  
1 row(s) retrieved.  
  
select count(*) from sst,tempmissingvalue where timeval between 726800 and 726900 and  
lat between 70 and 80 and sst <> missing_value;  
  
      (count(*))  
              172  
1 row(s) retrieved.  
  
Database closed.  
  
See?  
© Barrodale Computing Services Ltd. 2010  
Press the Enter key to continue ...
```

The VMware Player interface includes a taskbar at the bottom with a "Computer" icon, a "UFI Demo" window, and system tray icons for network, volume, and power. The system tray shows the date and time: "Mon Jul 19, 4:28 PM". A VMware logo is visible in the bottom right corner. A status bar at the very bottom reads "To direct input to this virtual machine, press Ctrl+G."



```
InformixVA-11.50.UC7-DE-SLES11-x86_32-v4.0.5 - VMware Player  File  VM  Help  ▾
UFI Demo
File  Edit  View  Terminal  Tabs  Help
UFI can provide a convenient means for just reading a file into an existing
(permanent) table, such as the following (this may take a few seconds):

dbaccess -e ufitest -

Database selected.

-- define a permanent table.
create table sst_historical_values(filename varchar(100),
                                timeval float,
                                lat smallfloat,
                                lon smallfloat,
                                sst smallfloat);

Table created.

insert into sst_historical_values select '/home/informix/Desktop/UFItest/netcdf/sst.m
ean.nc',
    * from sst where timeval between 726800 and 726900 and lat between 70 and 8
0;
2700 row(s) inserted.

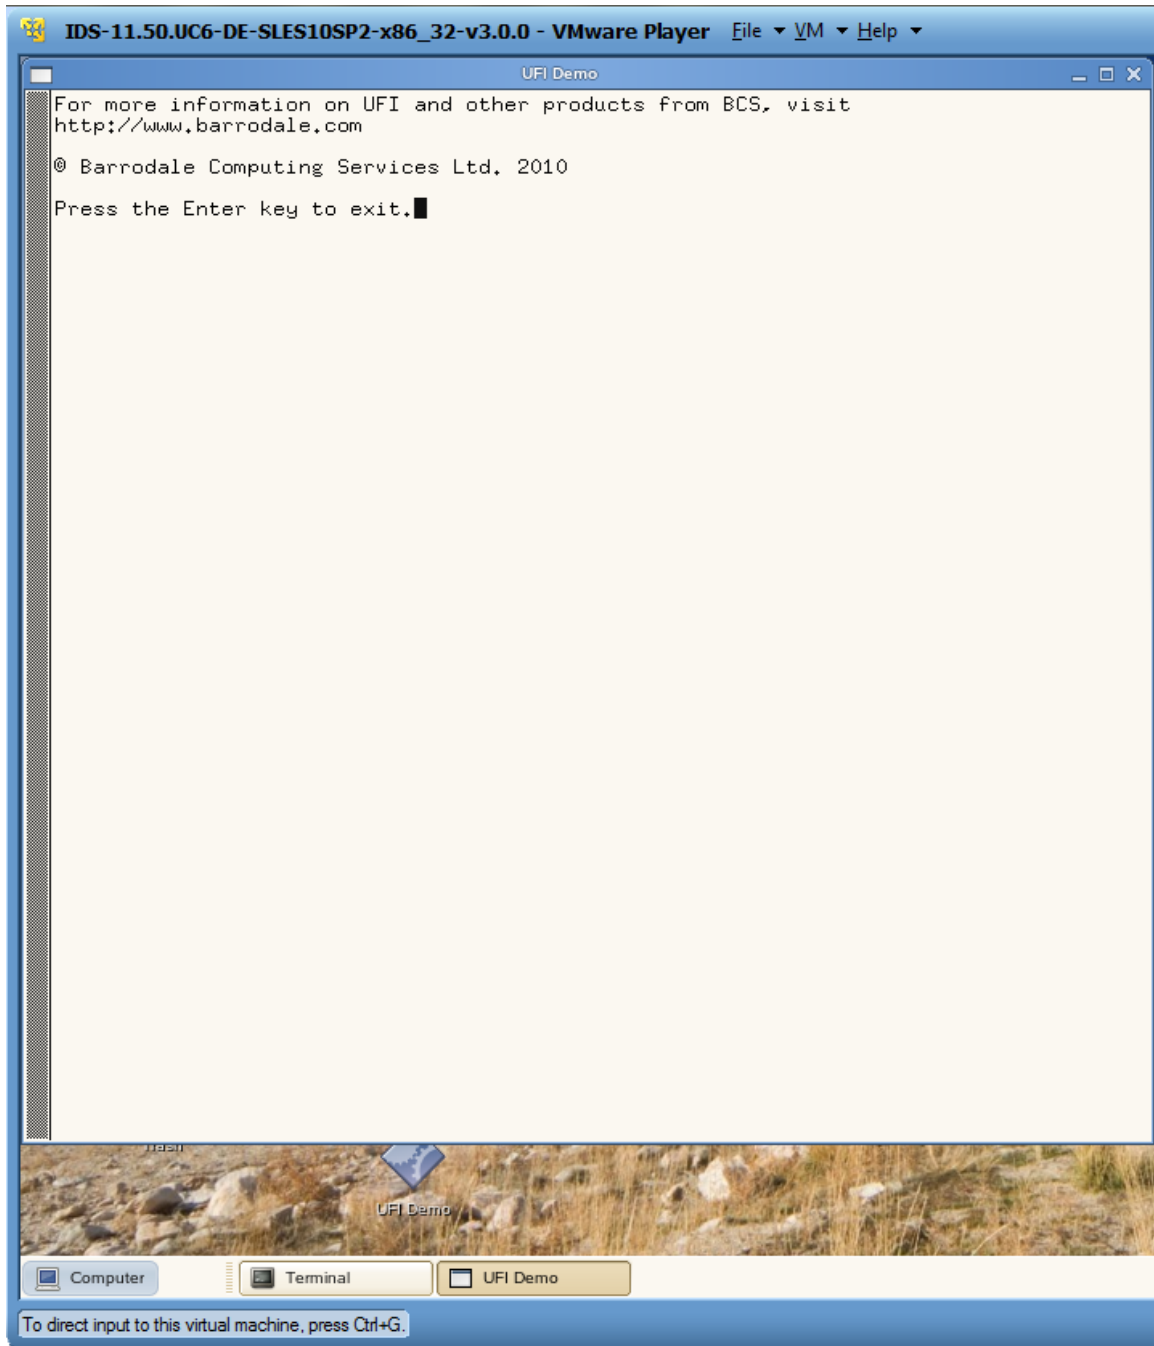
Database closed.

© Barrodale Computing Services Ltd. 2010

Press the Enter key to continue ...
```

Computer UFI Demo Mon Jul 19, 4:29 PM

To direct input to this virtual machine, press Ctrl+G. vmware



For More Information ...

To download a time-limited trial copy of UFI with the IDS Developer Edition Virtual Appliance, visit our website at <http://www.barrodale.com/bcs/bcs-ufi-doc-download>.

To discuss other trial options or to purchase UFI, please contact us at bcsInfo@barrodale.com.

To learn more about BCS and our other products, please visit our website at <http://www.barrodale.com>.