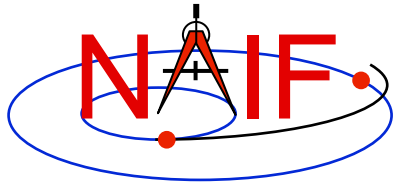


Navigation and Ancillary Information Facility

Introduction to Kernels

January 2012



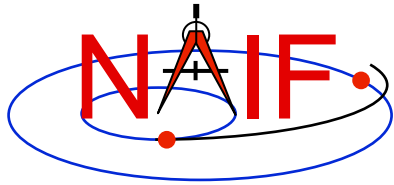
What is a SPICE “Kernel”

Navigation and Ancillary Information Facility

“Kernel” means file

“Kernel” means a file containing ancillary data


“Kernel” means a file containing "low level" ancillary data that may be used, along with other data and SPICE software, to determine higher level observation geometry parameters of use to scientists and engineers in planning and carrying out space missions, and analyzing data returned from missions.



SPICE Kernels Family

Navigation and Ancillary Information Facility

- **SPK**
 - Spacecraft and Planet Ephemeris
- **PcK**
 - Planetary Constants, for natural bodies
 - » Orientation
 - » Size and shape
- **IK**
 - Instrument
- **CK**
 - Pointing (“C-matrix”)
- **EK**
 - Events, up to three distinct components
 - » ESP: science plan
 - » ESQ: sequence
 - » ENB: experimenter’s notebook
- **FK**
 - Reference frame specifications
- **SCLK**
 - Spacecraft clock correlation data
- **LSK**
 - Leapseconds
- **Meta-Kernel** (a.k.a. “FURNSH kernel”)
 - Mechanism for aggregating and easily loading a collection of kernel files
- **DSK (under development)**
 - Digital shape kernel
 - » Tessellated plate model
 - » Digital elevation model



Text and Binary Kernels

Navigation and Ancillary Information Facility

SPICE **text** kernels are:

- text PCK (the most common type of PCK)
- IK
- FK
- LSK
- SCLK

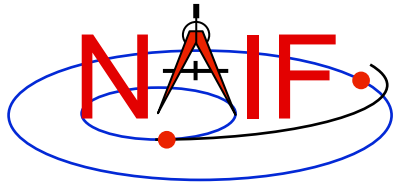
- MK (“Furnsh” meta-kernel)

SPICE **binary** kernels are:

- SPK
- binary PCK (exists only for Earth and moon)
- CK
- ESQ (part of the E-kernel)
- DBK (database kernel)

- DSK (digital shape kernel)*

* New kind of kernel under development




SPICE Kernel Forms

Navigation and Ancillary Information Facility

- **Binary form: SPK, binary PCK, CK, EK/ESQ¹, DSK**
 - Binary kernels are not human-readable and require the use of Toolkit software to examine the data contents.
- **Text form: text PCK, IK, FK, LSK, SCLK, FURNISH (MK)**
 - Text kernels contain only printing characters (ASCII values 32-126), i.e. human-readable text.
- **“Transfer” form of a binary kernel**
 - This is an ASCII representation of a binary kernel
 - Was used for porting the file between computers with incompatible binary representations (e.g. PC and UNIX)
 - Use of the transfer format is no longer needed for porting due to the run-time translation capability added to SPICE long ago
 - » But it is one way to convert a non-native binary kernel into native format, needed for modifying the kernel or improving read efficiency

[1] The ESP and ENB components of the EK might be binary, or text, or html, depending on specific implementation.



Text Kernel Data

Navigation and Ancillary Information Facility


- A text kernel contains data provided in simple assignments of the form:

`variable name = value(s)`

or

`variable name += value(s)`

- A text kernel can, and should, also contain descriptive meta-data that characterize the data



Example Text Kernel

Navigation and Ancillary Information Facility

This is a sample SPICE text kernel. The `\begindata` and `\begintext` markers on lines by themselves set off the start of data and text blocks respectively.

Note the forward slash!

```
KPL/<kernel type>
\begindata
  NAME           = 'Sample text value'
  NaMe          = 'Keywords are case sensitive'

  NUMBERS        = ( 10.123, +151.241, -1D14 )
  NUMBERS        += ( 1.0,    1,      -10    )
  NUMBERS        += ( 1.542E-12, 1.123125412 )

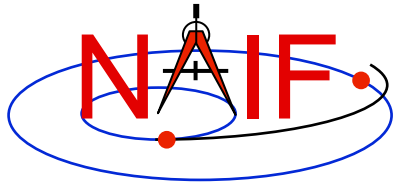
  NAIF_BODY_NAME += ( 'SPEEDO', 'NEETO' )
  NAIF_BODY_CODE += ( -678,    -679    )

  TIME           = @1972-JAN-1

\begintext
  < some comments about the data >
\begindata
  < more data, given again in keyword = value syntax >
\begintext
  < etc., etc. >
```

Note the backward slashes!


- Text kernels can contain characters, times, and numeric values.
- Text kernels can contain scalar or vector values.
- For more information see “Kernel Required Reading.”



Text Kernel Operators

Navigation and Ancillary Information Facility

- **An assignment using the “=” operator creates a variable name with one or more associated values.**
- **An assignment using the “+=” operator adds additional values to an existing variable. It creates a new variable if the referenced variable doesn’t already exist.**
- **The “@” symbol preceding a calendar date identifies a date string.**
 - **The string must not contain embedded blanks**
 - **The string will be parsed and converted to an internal double precision representation of that epoch.**
 - » **This conversion does not need a leapseconds kernel**
 - **The date is interpreted as ephemeris time (ET, also called TDB).**

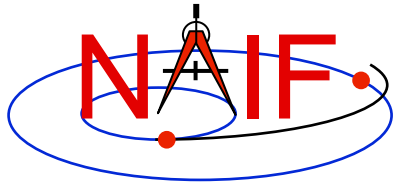


Example Binary Kernel

Navigation and Ancillary Information Facility

A binary kernel contains lots of non-printing (unintelligible) data, usually interspersed with occasional occurrences of ASCII characters.

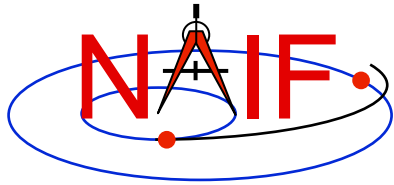
The only way to use a binary kernel is to read it, or add to it, using a SPICE routine or program.



Comments In SPICE Kernels

Navigation and Ancillary Information Facility


- **All SPICE kernels can and should contain comments—descriptive information about the data contained in the file.**
 - “Comments” are also known as “meta-data”
- **See the tutorial on comments for more information.**



Loading Kernels - 1

Navigation and Ancillary Information Facility


- To make kernels available to SPICE programs you “load” them.
- When you load a text kernel:
 - the file is opened
 - the kernel contents are read into memory
 - » variable names and associated values are stored in a data structure called the “kernel pool”
 - the file is closed
- When you load a binary kernel:
 - the file is opened
 - for SPK, CK, and binary PCK files, no data are read until a read request is made by Toolkit software
 - for ESQ files, the schema description is read, checked, and stored in memory at load time, but no data are read until a query/fetch is made
 - for all practical purposes the binary file remains open unless specifically unloaded by you



Loading Kernels - 2

Navigation and Ancillary Information Facility


- Use the FURNISH routine to load all kernels—text and binary.
 - Sample FORTRAN, C, IDL and MATLAB calls:
 - » `CALL FURNISH ('name.ext')`
 - » `furnsh_c ("name.ext");`
 - » `cspice_furnsh, 'name.ext'`
 - » `cspice_furnsh ('name.ext')`
- **Best practice: don't hard code filenames—list these in a “meta-kernel” and load the meta-kernel using FURNISH.**
 - `CALL FURNISH ('meta-kernel_name')` (Fortran example)
 - See the next page for more information on meta-kernels
- **Caution: "Transfer format" versions of binary kernels can not be loaded; they must first be converted to binary with the Toolkit utility program *tobin* or *spacit*.**



What is a “Meta-Kernel”

Navigation and Ancillary Information Facility

- **A meta-kernel is a file that lists names (and locations) of a collection of SPICE kernels that are to be used together in a SPICE-based application**
 - You can simply load the meta-kernel, causing all of the kernels listed in it to be loaded
- **Using a meta-kernel makes it easy to manage which SPICE files are loaded into your program**
- **A meta-kernel is implemented using the generic SPICE text kernel standards**
 - Refer to the Kernels Required Reading technical reference for details
- **The terms “meta-kernel” and “FURNSH kernel” are used synonymously**



Sample Meta-Kernel Contents (1)

Navigation and Ancillary Information Facility

- This is a sample meta-kernel that the Toolkit routine FURNISH could use to load a collection of kernels.


KPL/MK

```
\begindata
```

```
    KERNELS_TO_LOAD = (  
        '/home/mydir/kernels/lowest_priority.bsp',  
        '/home/mydir/kernels/next_priority.bsp',  
        '/home/mydir/kernels/highest_priority.bsp',  
        '/home/mydir/kernels/leapseconds.tls',  
        '/home/mydir/kernels/sclk.tsc',  
        '/home/mydir/kernels/c-kernel.bc',  
        '/home/mydir/kernels+',  
        '/custom/kernel_data/p_constants.tpc',  
    )
```

The commas
are optional

- The last file listed in this example (p_constants.tpc) demonstrates how to use the continuation character '+' to work around the 80 character limitation imposed on string sizes by the text kernel standards.



Sample Meta-Kernel Contents (2)

Navigation and Ancillary Information Facility

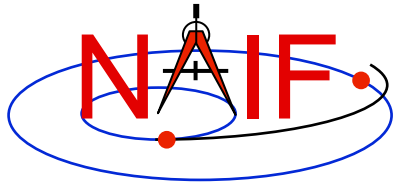
- This sample meta-kernel uses the `PATH_VALUES` and `PATH_SYMBOLS` keywords to specify the directory where the kernels are located.

```
KPL/MK
```

```
\begindata
```

```
  PATH_VALUES      = ( '/home/mydir/kernels' )
  PATH_SYMBOLS     = ( 'KERNELS' )
  KERNELS_TO_LOAD = (
    '$KERNELS/lowest_priority.bsp',
    '$KERNELS/next_priority.bsp',
    '$KERNELS/highest_priority.bsp',
    '$KERNELS/leapseconds.tls',
    '$KERNELS/sclk.tsc',
    '$KERNELS/c-kernel.bc',
    '$KERNELS/custom/kernel_data/p_constants.tpc'
  )
```

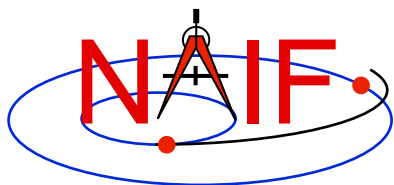
- Although the OS environment variable notation `$NAME` is used to refer to the symbols set by the `PATH_VALUES` and `PATH_SYMBOLS` keywords, **these symbols are NOT operating system environment variables and are set and used for substitution by SPICE only in the context of this particular meta-kernel.**
- The '+' continuation character described on the previous page may also be used to handle path value strings that exceed 80 characters.



Limits on Loaded Kernels

Navigation and Ancillary Information Facility

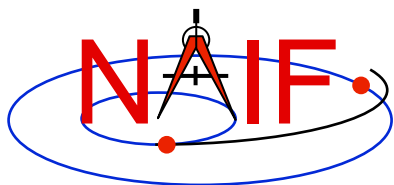
- **The number of binary kernels that may be loaded at any time is large, but limited.**
 - For SPK, CK, and binary PCK files:
 - » Loaded SPKs + Loaded CKs + Loaded binary PCKs ≤ 1000
 - For ESQ files:
 - » Loaded ESQs ≤ 20
 - For all kernels:
 - » Loaded kernels ≤ 1300
 - Assumes each has been loaded only once, and not unloaded.
- **There are also limits on the number of keywords and values for all loaded text kernels:**
 - Maximum number of keywords is 5003.
 - Maximum number of numeric data items is 200,000.
 - Maximum number of character data items is 4000.



Kernel Precedence Rule

Navigation and Ancillary Information Facility

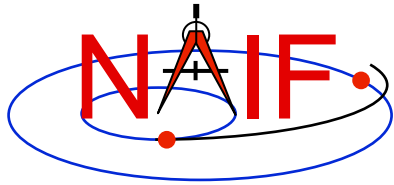
- **The order in which SPICE kernels are loaded at run-time determines their priority when requests for data are made**
 - For binary kernels, data from a higher priority file will be used in the case when two or more files contain data overlapping in time for a given object.
 - » For SPKs, CKs and binary PCKs the file loaded **last** takes precedence (has higher priority).
 - » Priority doesn't apply to ESQ files – all data from all loaded files are available.
 - If two (or more) text kernels assign value(s) to a single keyword using the “=” operator, the data value(s) associated with the last loaded occurrence of the keyword are used—all earlier values have been replaced with the last loaded value(s).
 - Orientation data from a binary PCK always supersedes orientation data (for the same object) obtained from a text PCK, no matter the order in which the kernels are loaded.



Unloading Kernels

Navigation and Ancillary Information Facility


- **The unloading of a kernel is infrequently needed for FORTRAN or CSPICE applications but is essential for Icy and Mice scripts**
 - Because of the way IDL and MATLAB interact with external shared object libraries any kernels loaded during an IDL or MATLAB session will stay loaded until the end of the session unless they are specifically unloaded
- **The routines KCLEAR and UNLOAD may be used to unload kernels containing data you wish to be no longer available to your program.**
 - KCLEAR unloads all kernels and clears the kernel pool
 - UNLOAD unloads specified kernels
 - KCLEAR and UNLOAD are only capable of unloading kernels that have been loaded with the routine FURNISH. They will not unload any files that have been loaded with older load routines such as SPKLEF (those used prior to availability of FURNISH).
- **Caution: unloading text kernels with UNLOAD will also remove any kernel pool data provided through the kernel pool API (P*POOL)**



Backup

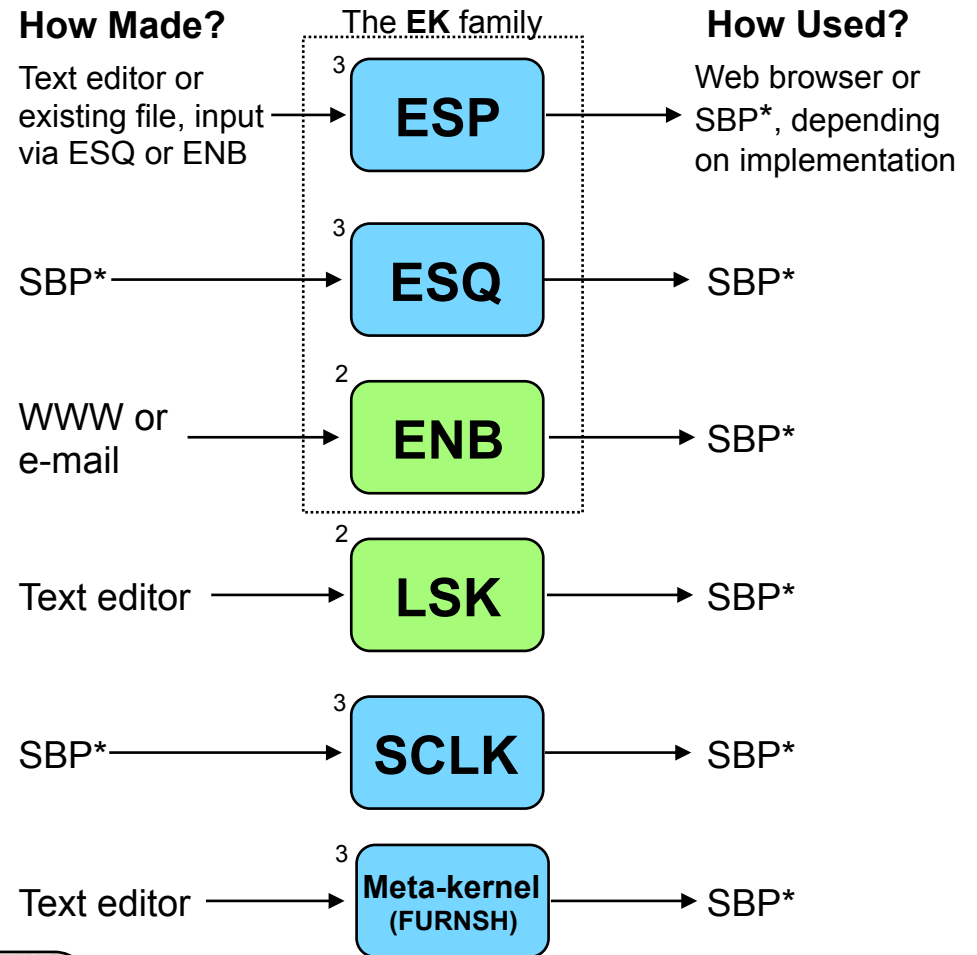
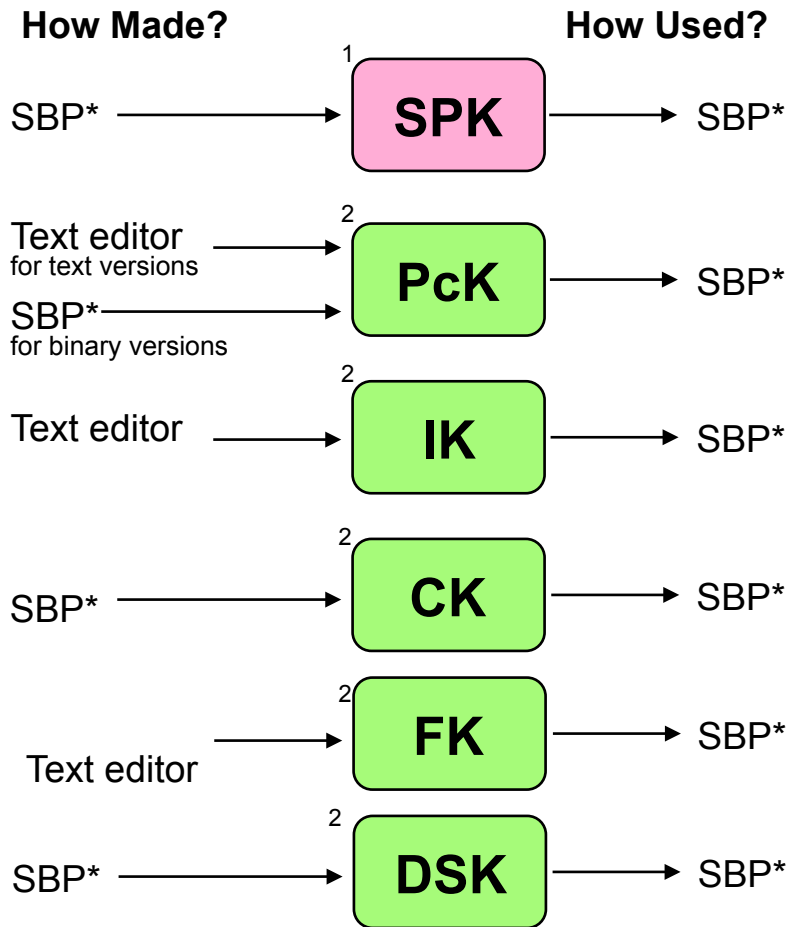
Navigation and Ancillary Information Facility

- **How kernels are made and used**
- **Why and how kernels are modified**
- **SPICE data structures hierarchy**
- **Problems making text kernels**



How Kernels are Made and Used at JPL

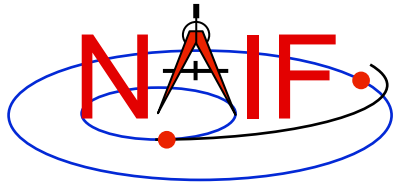
Navigation and Ancillary Information Facility



Who usually makes the kernels at JPL?

- 1 ● NAV and NAIF This represents current practice for most JPL missions, but is by no means a requirement. **Anyone can make SPICE files.**
- 2 ● NAIF
- 3 ● NAIF or other

***SBP** = SPICE-based program that uses modules from the SPICE Toolkit. In some cases the Toolkit contains such a program already built. In some cases NAIF may have such a ready-built program that is not in the SPICE Toolkit.



Why & How Kernels are “Modified” - 1

Navigation and Ancillary Information Facility

File Type

Why Modified

How Modified

SPK

- To add metadata (comments)
- To merge files or subset a file
- To correct/revise an object ID

- COMMNT, SPACIT or SPICELIB module
- SPKMERGE
- BSPIDMOD

PcK

Text version

- To revise data values
- To add additional data items and values

- Text editor
- Text editor

IK

- To revise data values
- To add additional data items and values

- Text editor
- Text editor

CK

- To add metadata (comments)
- To merge files
- To revise the interpolation interval
- To subset a file

- COMMNT, SPACIT, or SPICELIB module
- DAFCAT, CKSMRG
- CKSPANIT, CKSMRG
- CKSLICER

FK


- To revise data values
- To add additional data items and values

- Text editor
- Text editor

DSK

- To add metadata (comments)
- To merge files or subset a file


- COMMNT, SPACIT or SPICELIB module
- DSKMERGE



Why & How Kernels are “Modified” - 2

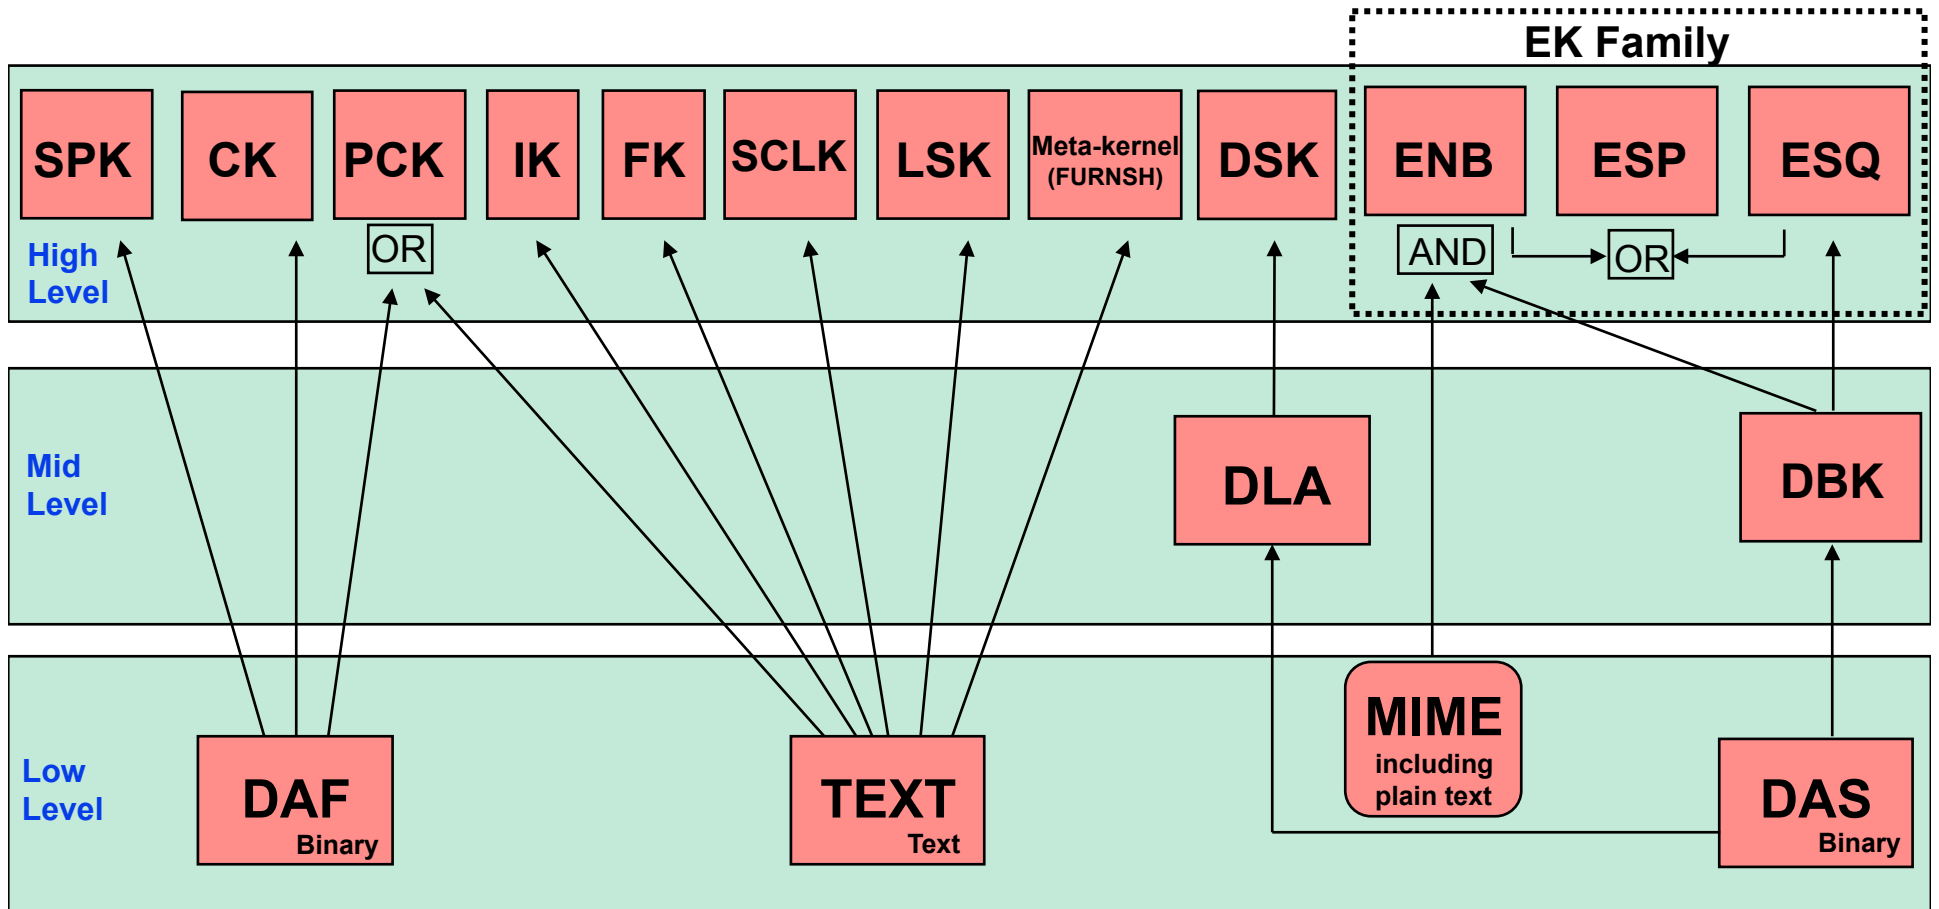
Navigation and Ancillary Information Facility

<u>File Type</u>	<u>Why Modified</u>	<u>How Modified</u>
The EK family		
ESP	<ul style="list-style-type: none"> -To add, revise or delete “data” -To add metadata (comments) 	<ul style="list-style-type: none"> - (Depends on implementation) - (Depends on implementation)
ESQ	<ul style="list-style-type: none"> -To add additional data -To revise data -To delete data -To add metadata (comments) -To merge files 	<ul style="list-style-type: none"> - Toolkit modules - Toolkit modules - Toolkit modules - COMMNT, SPACIT or SPICELIB module - (under development)
ENB	<ul style="list-style-type: none"> -To change entry status (public <--> private) -To delete an entry 	<ul style="list-style-type: none"> - WWW - WWW
LSK	<ul style="list-style-type: none"> - To add a new leapsecond 	<ul style="list-style-type: none"> - Text editor
SCLK	<ul style="list-style-type: none"> - To add metadata 	<ul style="list-style-type: none"> - Text editor
Meta-kernel (FURNISH)	<ul style="list-style-type: none"> - To revise contents in any way 	<ul style="list-style-type: none"> - Text editor



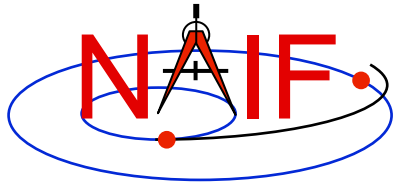
SPICE Data Structures Hierarchy

Navigation and Ancillary Information Facility



DAF = Double Precision Array File DSK = Digital Shape Kernel (under development)
 DBK = Data Base Kernel DLA = DAS Linked Array (under development)
 DAS = Direct Access, Segregated

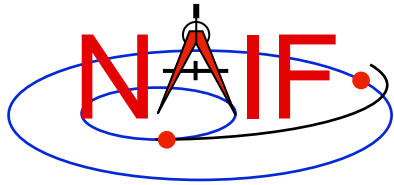
Excepting MIME, each of these data structures is built entirely of SPICE components.
 Pck files are usually text-based, but binary versions exist for the earth and moon. The ESP has been implemented using both the ENB and ESQ mechanisms. The DBK is a SQL-like, homebrew database.



Problems Making Text Kernels

Navigation and Ancillary Information Facility

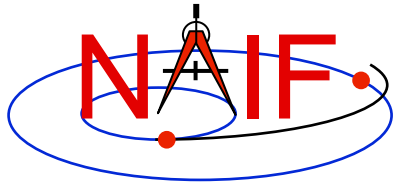
- **Cutting/pasting complete, or pieces of, data assignments or `\begindata` or `\begintext` markers into a text kernel can cause a problem**
 - It may result in insertion of non-printing characters or incorrect end-of-line terminations
 - This is not a problem for comments, but it is probably best to treat all portions of a text kernel the same
- **If creating a text kernel by editing an existing one:**
 - first save a backup copy
 - be sure you are starting with a file in native format for the computer you are using: either Unix/Linux/Mac or Windows
 - be sure to insert a final end-of-line marker at the end of your last line of data or text
 - » Press “return”



Some Useful Tools - 1

Navigation and Ancillary Information Facility

- **For a Unix or Linux (including Mac) environment**
 - To display all non-printing characters, tab characters as “^I”, and place a “\$” character at the end of each line
 - » `cat -et <file name>`
 - » **How do end-of-line markers appear when displayed in a text file using the `cat -et` command?**
 - Unix/Linux/Mac: \$ (line feed)
 - Windows: ^M\$ (carriage return followed by line feed)
 - Display the file type, language used, and end of line marker
 - » `file <file name>`
 - » **Examples using Unix and Windows (“PC”) versions of the SPICE leapseconds kernel:**
 - `file naif0010.tls`
`naif0010.tls: ASCII English text`
 - `file naif0010.tls.pc`
`naif0010.tls.pc: ASCII English text, with CRLF line terminators`



Some Useful Tools - 2

Navigation and Ancillary Information Facility

- **For a Unix or Linux (including Mac) environment**
 - To convert a Unix/Linux/Mac text kernel to Windows (“DOS”) style:
 - » `unix2dos <filename>`
 - To convert a Windows (“DOS”) style text kernel to Unix/Linux/Mac style:
 - » `dos2unix <filename>`
 - For converting either style text kernel to the other style, use the SPICE “bingo” executable
 - » The bingo program and User Guide are available only from the NAIF/Toolkit/Utilities web page:
 - <http://naif.jpl.nasa.gov/naif/utilities.html>
- **More information**

In Wikipedia, search on “newline” or “unix2dos”