

Fermi National Accelerator Laboratory

**FN-442
2320.000**

**Organizing, Maintaining, and Distributing
Software Products***

**P. Heinicke, T. Nicinski, P. Constanta-Fanourakis, D. Petravick, R. Pordes,
D. Ritchie, and V. White
Fermi National Accelerator Laboratory
P.O. Box 500, Batavia, Illinois 60510**

October 1986

***Submitted to the Proceedings of the Digital Equipment Users Society, San Francisco, CA, October 6-10, 1986.**



Operated by Universities Research Association Inc. under contract with the United States Department of Energy

Organizing, Maintaining, and Distributing
Software Products

Peter Heinicke, Tom Nicinski,
Penelope Constanta-Fanourakis, Donald Petravick,
Ruth Pordes, David Ritchie, Vicky White

Fermi National Accelerator Laboratory
Computing Department / MS120
P.O. Box 500, Batavia, IL 60510

ABSTRACT

The Computing Department at Fermilab develops and maintains software used at more than 30 different sites. A general methodology has been devised to keep track of and distribute the software at these different sites. Experience over the past year has proven the usefulness and efficacy of the method.

1 Introduction

The Fermi National Accelerator Laboratory (Fermilab) is a facility dedicated to basic research in the field of high energy physics. This research takes the form of "experiments", which are conducted by groups of physicists. The experiments are highly computerized; there are usually one or more minicomputers devoted to the tasks of data acquisition and analysis of the experimental data.

Most experiments have at least one VAX or MicroVax computer, as well as one or more PDP-11 computers, and possibly various programmable microprocessors. There are many different experiments either actively taking data or preparing to do so at any one time.

The Data Acquisition Software Group of the Fermilab Computing Department provides software support for the experiments. The Data Acquisition Software Group's role is to develop a wide range of useful software for data acquisition and analysis. Experimenters use the software to perform the required online data acquisition and analysis for their experiment. In some cases, the software is used in a turnkey manner; more often, however, it is used as the basis for more elaborate and

experiment-specific software. In the latter case, the experimenters obtain the basic package and then do their own software development to customize it to their particular needs.

Software is targetted for PDP-11 or VAX computers; the target operating system environment is RT-11/RSX-11M or VMS. Other targets are microprocessors, such as 68020's, etc. These target computers are located at approximately 30 different sites scattered over the 6800 acres of Fermilab. The VAX's and MicroVAX's at these sites are connected to one another via DECnet. These VAX's (or the Central Facility VAX Cluster) are used by the experimenters for software development in enhancing the supplied software as well as for online data acquisition and analysis. Software is transferred to these machines via DECnet from the Data Acquisition Software Group's Development VAX. It is also transferred via magnetic media to the computers not connected via DECnet. These include the PDP-11's (not connected mainly due to memory limitations) and the microprocessors.

Additionally the software may be transferred to the collaborating universities and research institutions which participate in Fermilab experiments. This transfer occurs so that the experimenter may continue software development activities for a Fermilab experiment while residing at the home institution or so that the experimenter may test apparatus under construction with components of the software intended for the experiment.

With so many sites and so much software in use at these sites, we quickly realized that some systemization of the task of organizing, maintaining, and distributing the software was mandatory. Keeping track of the software at the various sites, although a formidable job, is nevertheless a necessary one--we must be able to offer assistance with the current version of the software at hand.

A requirement on the systemization was that it must support having different versions of the same software at different sites or even at the same site.

While it might be possible in principle to arrange the same version of the software at all sites, in practice it does not occur. One of the most important reasons is that an ongoing experiment does

not necessarily want to avail itself of the latest enhanced version of a piece of software; bugs or side effects may be introduced which might complicate the primary task of monitoring the experiment. Even when an experiment decides that the new features outweigh any risks of complication, it is extremely important that the experiment be able to switch back to the previous version as quickly as possible. The motivation may be to retreat from a software enhancement because it itself was found to have problems or because one wishes to rule out software changes as a cause of changes in the data being monitored. When the latter occurs, one then wishes to go forward again to the latest version.

This paper describes how we have organized our software development and support efforts to satisfy these considerations. In what follows, we describe the organization of our software into "Products", how these Products are created, maintained and versioned, and how this Product organization is used in the distribution of software to the target VAX computers, and from there to other target computers when necessary.

1.1 What Is A Product?

A Product is an arbitrary group of logically connected directories and files (stored on a VAX/VMS system) and referred to by a Product name and optionally by qualifying names, such as the Version number, target operating system, or hardware interface. The Product name is a printable ASCII string describing the group in a mnemonic way. For each Product name there is a single development version of the product and/or one or more distribution versions. It is not necessary that a Product be developed by the Computing Department to participate in this scheme. However, the Product (the directories and files which comprise it) must be organized in a prescribed way. The constraints are relatively minor because we wanted the ability to include all kinds of software as products--not just those developed at Fermilab.

An example of a non-Fermi Product is KERMIT, a communications package. KERMIT VMS is the Product name for the VMS version of KERMIT.

When the source code contained in the development version of a Product is updated, either for maintenance or enhancement reasons, a new "Version" of the Product is generated. This may occur even if the source code of the Product is unchanged. For example, if a Product is rebuilt using new "versions" of code on which it depends (such as an object library), but which is not a part of the Product itself, a new version of the Product is still generated. A Product version is used to inform the user, developer, and Product maintainer of not only which level of source code of the product it contains but also the entire state of the Product, its dependencies on other software Products, etc.

As a simple example of a Product with different Product versions, consider the KERMIT product for RT-11, where each version reflects the update level and the language.

Version V1.0 of the KERMIT RT Product refers to the first Pascal version of RT-11 KERMIT. When the MACRO-11 version became available, the users needed to decide whether to keep supporting the Pascal version. If they had, they could have renamed it to be the "KERMIT PASCAL" product (and call the other version the KERMIT MACRO RT product V1.0). Otherwise, they could have chosen to supersede it with V1.1 of KERMIT RT.

Products come in two flavors: "simple" and "compound." A simple Product consists of a collection of software which is expected to be used, upgraded to a new version, and distributed to target sites independent of the state of other software Products. The decision to organize a product as a "simple" one is basically that of the developer; it is a statement that this Product is somehow basic and not further made up of Products.

This does not necessarily mean that the Product was not dependent upon other software external to the Product when it was "built" (compiled, linked, etc.). Nor does it necessarily mean that the Product requires no other software Product in order to function.

For example, many of our Products are written in FORTRAN. These are definitely dependent upon the FORTRAN compiler and the FORTRAN Run Time Library--both of which are external to the product and which (in the case of the Run Time Library, at

least) are required in order for the Product to function.

A compound Product is a collection of different "component" Products (either simple or compound), frequently used together. These Products do not necessarily have to be dependent upon each other although in many cases they are. They may be grouped together only for ease of distribution of many small Products which change infrequently. Alternatively, they may be grouped together because of dependencies on each other; hence, a change in a component Product would indicate that a new version of one or more of the other components is either necessary or desirable. DEC's ALL-IN-1 system is an example of something that is structurally similar to a compound Product.

1.2 Goals

The Data Acquisition Group is primarily responsible for designing, developing, and maintaining software as well as supporting the end users of the software. The distribution and installation of the software is only a peripheral activity. To permit us to spend more time on software development, we have devised a Product Specification and specialised procedures, whose goals are:

- o Provide a Uniform Product Specification.

The Product specification is meant to provide system management tools and the user with a uniform interface to the software we are responsible for. The specification includes

- o the directory structure of the files in a Product (defined to be a tree structure)
- o a list of required and optional files,
- o the naming conventions for these files and directories,
- o how logical names should be used.

- o Keeping Track of Product Versions on a System.

Different sites use different versions of a Product creating a need to maintain a database of which Products and versions reside on a particular system. This functionality is provided by a system management tool we call

SITE_PRODUCTS.

o Simplification of Product Distribution.

We need to automate the distribution of versions of Products to remote sites (making use of DECnet) and the installation of the Products on the target site. Such automated procedures are needed both for efficient use of our time and to minimise the risk of errors or omissions.

o Transportability to External Sites.

Although restrictions are placed on a Products structure and interaction with users (how the Product is distributed and how the system manager treats it), it is still necessary to permit the Product to be easily installed and used on systems which do not follow our methodology.

o Permit Switching Between Product Versions.

In order to maintain and improve existing Products, and have the new releases accepted by experimenters, there is a need to allow the use of the latest version of a Product, but also to instantly and transparently "switch" to using a previous version residing on the same system.

The ability to switch between versions on the same system is also important for Product developers and maintainers. A user may discover a bug at a previous release of the Product - and the Product maintainer is then able to check for the bug in that release just by switching to it. This capability is provided by PRODUCT_SETUP and the database of Products and their versions (maintained by SITE_PRODUCTS).

o Permit the Composition of a Product to be Known Precisely

We make extensive use of DEC CMS (Code Management System) and MMS (Module Management System) to control the source code release level of a Product and to automate the construction of that product from its sources and any other libraries etc. it may be dependent on. However in situations where a Product may be dependent on libraries in other

Products - the specific version of the library-related Products used must be both controllable and forever known. The time-stamps of the individual files as used by MMS are not sufficient to control such inter-dependencies.

The procedures which we call BUILD permit the dependencies of one Product on another, either as a part of a compound Product, or just as a required but separate piece of software, which must be present in order to build the Product, to be expressed in a formal way. From this formal specification the order of creation of the component parts can be determined and the business of creating a very large software Product can be automated in a foolproof way.

1.3 The Results

All the management tools we have developed are written as DCL command procedures. DCL command procedures were chosen because of speed of implementation, and because we underestimated the full extent of the project we were undertaking.

The remainder of this paper will discuss the concepts and management tools introduced above which together allow us to achieve the goals outlined in the previous section. These include: Specification of a Product, use of the BUILD procedures, the SITE_PRODUCTS, DISTRIBUTE and PRODUCT_SETUP procedures.

2 Specification Of A Product

The Product Specification provides system management tools and the user with a uniform interface to the software. We have written a 50-page specification of a Product including the mandatory and recommended requirements thereon. The Product Specification addresses three areas:

- o Directory tree structure and the files in a Product.
- o Logical names to be defined (associated with the Product).
- o Required and optional command procedures and how they are used. (definition of parameters).

2.1 Directory Tree Structures

Products reside under rooted directories. Actually, two rooted directories are associated with a particular Product. The "Version" Root is the rooted directory for a particular version of a Product. This is the rooted directory that a user will see when using a Product. Version Rooted directories reside under an "Umbrella" Rooted directory. The Umbrella Rooted directory contains all the versions of a Product. However, more than one Product and its versions can reside under the same Umbrella Root. For example:

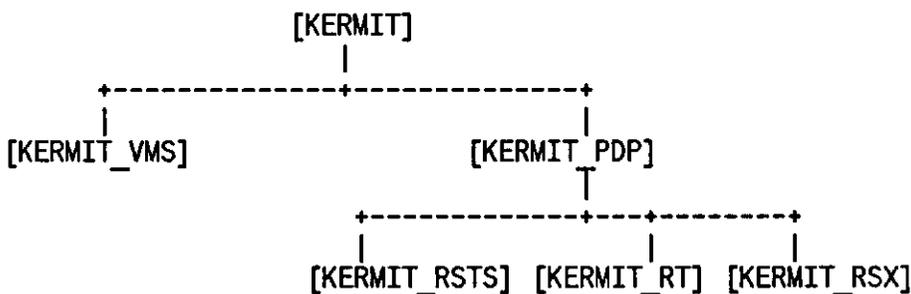


Figure 3-1

The leaves are products, while [KERMIT] is the Umbrella Root. The Version Roots for the Products are [KERMIT_VMS], [KERMIT_RSTS], [KERMIT_RT], and [KERMIT_RSX]. The directory [KERMIT_PDP] is an intermediate directory, which could also represent a Product, or an Umbrella Root (the interpretation is up to the Product developer).

For each Product version, there is a set of required and optional directories:

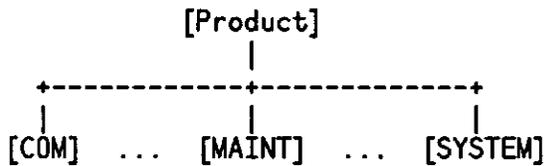


Figure 3-2

The [PRODUCT] directory is the Version Rooted directory, while [COM] and [SYSTEM] are required, and [MAINT] is an optional directory. Beyond these directories, the developer can use any tree structure (under the Product's version rooted directory).

2.2 Logical Names

To keep Products site-independent, logical names are used to point to different files. All logical names should be defined in terms of one logical name:

```
'product'$ROOT
```

which is the rooted logical name pointing to the PRODUCT's Version Root. By changing 'product'\$ROOT's definition (with PRODUCT_SETUP), a user can easily "switch" between different versions of a Product. In the example in Figure 3-1, the rooted logical name for HORSE is

```
$ SHOW LOGICAL KERMIT_RT$ROOT
"KERMIT_RT$ROOT" ≡ "disk:[KERMIT.
                    KERMIT_PDP.KERMIT_RT]"
```

2.2.1 Logical Name Tables

Some Products require a large number of logical names to be defined. Ideally, users should only see logical names that they require, which implies that they should be process logical names. But, defining many logicals can be quite time consuming. The solution to make the logical names system wide was rejected for aesthetic and performance reasons, in that the system logical name table (LNM\$SYSTEM_TABLE) would become cluttered as the number of Products grew.

Instead, shareable (system wide) logical name tables are created for each Product. When a Product is started up, it defines its logical names in the table created for it. This makes the logical name tables (and the logical names) invisible unless they are required.

Keeping a Product's logical names within its own logical name table keeps the system clean and allows for easy switching between logical names defined for different versions of a Product. It also helps when looking for all the logicals associated with a particular Product when you are on a large system with many logicals defined.

To use a Product, a user invokes `PRODUCT_SETUP` (described later) to "link" the logical name table into his/her logical name table search list.

2.3 Required Files

The Product specification requires that each product provide two command files, of defined names, to be implicitly invoked at system bootstrap time and when a user wants to use the Product. All products must provide these files in a particular directory for the Product version. The specification also recommends a Help file to be provided with each Product; this is automatically included in the general Product Help library when the Product is entered into the `SITE_PRODUCTS` database.

`[COM]SETUP.COM` is used to define logical names and symbols on a per process basis. That is, the user invokes `SETUP.COM` (normally at login time) if there is a need to use the Product.

`[SYSTEM]PRSTARTUP.COM` is used during system boot time (Product startup) to define shareable logical names in the logical name table generated for the Product, and to perform any other operations which affect the Product system wide (such as `INSTALLING` files, loading device drivers, starting a queue, etc.) and other privileged initialization functions.

3 BUILD And Developing The Products

The BUILD procedure is used to construct a Product based upon its dependencies on other Products. BUILD takes into account that a Product may:

- o Depend on other Products.
- o Depend on specific versions of other Products.
- o Incorporate other Products totally within it.

The construction of a Product consists of compiling and linking the software comprising the Product.

A Product developer uses a Product Maintenance Language (PML) file to describe how a product is dependent upon other products. Only the immediate dependencies need to be described, since BUILD recursively uses the dependent Product's PML files to generate a final list (a Product Maintenance Output (PMO) file) which sequentially describes the order in which Products should be built (to satisfy all dependencies).

For example, the product KERMIT_VMS is to be "BUILT":

- o KERMIT_VMS is dependent upon an another product called GET_PORT
- o KERMIT_PDP is dependent upon KERMIT_RT, KERMIT_RSX, and KERMIT_RSTS.

BUILD would determine that the Products would need to be built in the following order:

```
GET PORT
KERMIT_VMS
KERMIT_RT
KERMIT_RSX
KERMIT_RSTS
KERMIT_PDP
```

BUILD then will construct the Products in the appropriate order to generate the final Product. To save time, BUILD will not construct a Product if the required version already exists.

The actual details of construction of each of the component pieces are left up to the component piece of software. We normally use DEC CMS and MMS wherever possible. This is especially useful in conjunction with our methodology of one development version of a Product and multiple distribution versions. By having a single CMS library in the development version of each Product and creating classes for each source release level we avoid the need to keep the sources with or for each version of the Product. We can always recreate any version at any time. This saves disk space and also provides a centralized record of who changed the software and when.

4 SITE_PRODUCTS - System Management Of Products

SITE_PRODUCTS was developed to keep track of which versions of which Products reside on a system. It not only maintains a database of Products and their versions, but it schedules the starting up of Products at system boot time (or any other time) and the shutting down of Products. SITE_PRODUCTS avoids the need for the system manager to change the system specific startup command procedure (SYSTARTUP) every time a Product or a version of a Product is added, modified, or removed.

Products are made "known" to SITE_PRODUCTS (this should not be confused with the known files of the VAX/VMS INSTALL Utility). The "Known Product List" file, maintains this information.

For each known product, SITE_PRODUCTS maintains a "Product Version List" file which resides under the product Umbrella directory. The Product developer is able to add, modify, and remove Product versions without requiring privileges (only access to the particular Product's area is required).

The SITE_PRODUCTS procedures point to the Known Product List using a logical name. Users can use SITE_PRODUCTS to maintain their own Known Product List, and Product Version Lists. This can be extended for use on a VAX Cluster system, where a common Known Product List is used to startup (shutdown) all Products common to all nodes in the Cluster. Then, by redefining the logical name, a

node-specific Known Product List can be used to manipulate software Products licensed (or useable) only for that particular machine.

SITE_PRODUCTS allows the addition, modification, and removal of Products and Versions. These operations only modify the Known Product List and Product Version Lists, not the actual files of the Products. When a Product version is declared to be the default version on a system, its Help file is included in a general Product Help library (if one exists) and also a Bulletin is posted on the system (if the Bulletin Product is available).

For each Product, the Known Product List maintains the Product's name, the specification of the Umbrella Root, and other miscellaneous information. Associated with each Product version in the Product Version List is a directory path from the Umbrella Root to the rooted directory for the Product version.

When a Product is started up by SITE_PRODUCTS, a shareable (system wide) logical name table is created to contain logical names defined by the Product. Then the Product specific startup command procedure is invoked. This procedure usually defines logical names, device drivers, starts up queues, installs privileged images, etc.

5 PRODUCT_SETUP - Using The Products

The final stage of any Product is its use. PRODUCT_SETUP is used to "setup" a product for use by a user. It also allows a user to choose which version of a product to setup. Setting up a Product involves the definition of logical names and symbols required for using the Product.

A symbol by the name of SETUP is used on all systems to invoke PRODUCT_SETUP. Users of a software Product such as our example KERMIT_VMS simply type

```
SETUP KERMIT_VMS
```

to use the default version of the Product and all its component sub-Products.

The ability to switch transparently between Product versions is provided by the logical name tables created for the Product. When switching between Product versions, `PRODUCT_SETUP` creates a new logical name table (which overrides the old table) and defines the logical names for that particular version. Therefore, different Product versions are not required to use the same logical names.

6 DISTRIBUTE - Distributing The Products

DISTRIBUTE provides a system manager on a remote machine the ability to copy Products, from an "Archive machine", and install them. Most of the time, DISTRIBUTE is used over DECnet, but it also provides a tape mode, which permits Products to be distributed and installed at external sites using Magnetic tape as a transfer medium.

DISTRIBUTE interactively queries the user for the information it needs. The questions are self explanatory, so that no documentation is normally required in order to obtain a Product. Besides the Product name and version, DISTRIBUTE asks where the Product should be placed (the disk and Umbrella Root), and whether the Product and its version should be declared to `SITE_PRODUCTS`.

When a Product is selected by the user, DISTRIBUTE uses that Product's Product Maintenance Output (PMO) file (generated during a BUILD) to determine which Component Products need to be copied over as part of the chosen Product. This provides all sites with a complete and consistent view of a Product. Products which are not constructed with BUILD and therefore have no PMO file can also be distributed - all files in the directory tree stemming from the Product version rooted directory will be taken to comprise the Product version.

DISTRIBUTE uses BACKUP save sets compatible with the VMSINSTAL utility (part of VAX/VMS). Because the Product conforms to the Product Specification, only one KITINSTAL file (used by VMSINSTAL) needs to be written for all Products. This frees the Product developer from writing code used strictly for the purpose of installing a product.

A complete log of software distributed, date, version and to where is maintained on the Archive machine.

7 Conclusions

The organisation of products and the procedures described in this paper have been in use for more than a year now. Hundreds of Products have been distributed to target sites. The sacrosanct nature of a Product version once built has enforced a strict discipline on program development and aided immensely in tracking down complicated problems where any one of a number of hardware and software variables could have been at the root of the problem. The procedures described were first developed for software to be executed on a VAX(VMS). We have found them such a useful aid for distribution, maintenance and archiving that we extended the concepts to cover software for other operating systems in use.

We have found the Standard Product specification to be extremely useful. Not only has it enabled us to write the management tools described but it has also helped enormously in the ease of understanding, maintaining and supporting our software. New members of the group and new users new to Fermilab can very quickly produce software to conform to the general specifications and obtain and use software that is available. It is much easier for any member of the group, regardless of particular area of expertise to be able to distribute, demonstrate, find bugs in, create a new version of any Product. New software Products produced elsewhere at Fermilab or at other institutions or vendors can be quickly added to the set of available software and made available in the same uniform way to all the users on site (via the same SETUP command). We package all software according to our minimum standards - give it a Product name, a version, keep all versions under a single Umbrella directory and define all logicals relative to a single root logical name pointing to the specific Product version. Following software Product "standards" has saved manpower also in enabling us to write general procedures. For example, the arrival of Microvaxes with limited disc space created a need to trim Products. A general

procedure which omitted all list, map and documentation files from a distribution version could be written because of the standards imposed, thus solving the problem in general for all software which we maintain or distribute.

This entire program of work was undertaken without a proper realization of the size of it - really as a non-serious sideline, which people did a little work on when the need arose. If we were doing it again we would better understand the benefits and scope of the project and would take it further than we have today. The database maintained by SITE PRODUCTS would be made extensible and easily accessible as a database. Some of the system management procedures would have been written in a high level language instead of DCL, thus increasing both their speed and extensibility.

8 Acknowledgements

Contributions to the ideas, definitions and procedures have been made at various times by all members of the Data Acquisition Software and DEC Systems Group in the Computing Department at Fermilab - which consist of the authors, David Berg, Eileen Berman, Andy Cohen, Terry Dorries, Arkady Lubinsky, Carmenita Moore, Liz Quigg, Dave Ritchie, Chip Kaliher, Nancy Hughart and Steve Kalisz. We also acknowledge helpful feedback from various users of the system (DISTRIBUTE in particular) ranging from on-site local system managers to experiment participants distributing software over DECnet from Italy.

9 References

PN's refer to Fermilab "Programming Notes"; IN's refer to Internal Notes. Documentation is available from the Computing Department Program Librarian.

IN 140 SITE PRODUCTS / Maintaining Known Products
by Tom Nicinski.

- IN 141 BULLETIN / Maintaining an Electronic Bulletin Board by Tom Nicinski.
- IN 157 Data Acquisition Software Group Product Specifications
edited by Ruth Pordes.
- PN 259 BUILD Procedure for Product Distribution by Penelope Constanta-Fanourakis.
- PN 261 Backup / Distribute Procedure for Product Distribution by Peter Heinicke.
- PN 262 Using VMSINSTAL with User-written Applications courtesy of Richard Aurbach of Monsanto, via DECUS.
- PN 269 PRODUCT SETUP User's Guide / Setting Up Products by Tom Nicinski.