

A Multiple Node Software Development Environment

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

Experimenters on over 30 DECnet nodes at Fermilab use software developed, distributed, and maintained by the Data Acquisition Software Group. A general methodology and set of tools have been developed to distribute, use and manage the software on different sites. The methodology and tools are of interest to any group developing and using software on multiple nodes.

Introduction

The Fermi National Accelerator Laboratory (Fermilab) is a facility dedicated to basic research in the field of high energy physics. 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, one or more PDP-11 computers, and possibly a few programmable microprocessors. Many different experiments either actively take data or prepare to do so simultaneously.

The Data Acquisition Software Group of the Fermilab Computing Department provides software and support for the experiments. Experimenters use the provided software to perform online data acquisition and analysis required for their experiment. In some cases, the software is used in a turnkey manner; more often, 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 customize it to their particular needs through their own software development efforts.

Software is usually targeted for PDP-11 or VAX computers. Other targets include microprocessors, such as 68020's, etc. Target computers are physically 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; (PDP-11's not connected mainly due to memory limitations and microprocessors).

Additionally, the software sometimes needs to be transferred to the collaborating universities and research institutions which participate in Fermilab experiments. The experimenter may then continue software development or equipment testing activities while physically located at the home institution.

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 is a formidable and necessary job. 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 to have the same version of the software at all sites, in practice this does not occur. 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 and reliably as possible.

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.

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 fit into this scheme. However, the Product (the directories and files which

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

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 COURIER product for VAXONLINE. Version V1.0 of the COURIER Product refers to the first released version of COURIER for VMS. It will normally have a product directory by the name of COURIER_V1_0. Later versions will have similar product directory names, e.g. COURIER_V1_2. If a UNIX version is developed, the support group would need to decide whether to keep the old name. If they decide to, they can rename the two products to be COURIER_VMS and COURIER_UNIX, or leave it as COURIER and COURIER_UNIX.

Products can be divided into two levels of complexity: "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 component 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.

Goals

The distribution and installation of the software is only a peripheral (but time consuming) activity. To permit us to spend more time on software development, we have devised a formal specification for "Products" and specialized procedures, whose goals are:

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 tree structure of the files in a product
- o a list of required and optional files,
- o the naming conventions for these files and directories,
- o how logical names should be used.

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.

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 minimize the risk of errors or omissions.

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.

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).

Permit the Composition of a Product to be Known Precisely. We make extensive use of DEC's CMS (Code Management System) and MMS (Module Management System) to control the source code version of a product and to automate the construction of that product from its sources and any other libraries etc. it may be dependent on. (CMS and MMS are similar to the SCCS and MAKE utilities on UNIX). 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.

The Resulting Tools

All the management tools we have developed are written as DCL command procedures. Any language could have been chosen, and the system could have been implemented in a more operating system independent way. DCL command procedures were chosen because of speed of implementation, and because we underestimated the full extent of the project we were undertaking. Future implementations of this functionality will probably use a different tool than DCL, since DCL is so slow, and unmodular. We are considering reimplementing the system in FORTRAN or an inexpensive 4GL.

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.

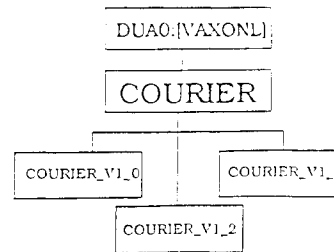
Specification of a Product

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).

Directory Tree Structures

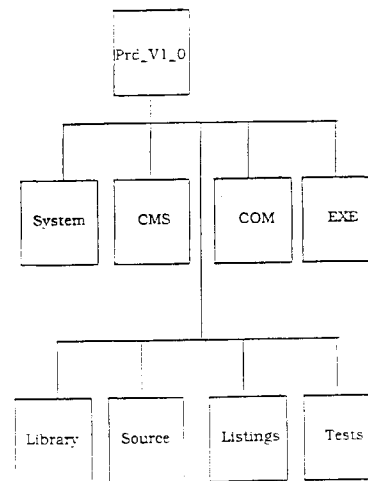
Products reside under rooted directories. Actually, two rooted directories are associated with a particular product. The "Version" Root is the rooted directory for a particular version of a product. This is the rooted directory that a user will see when using a product. Version Rooted directories reside 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:



Courier Product Directory

The leaves are products, while [VAXONL.COURIER] is the Umbrella Root. The Version Roots for the products are [COURIER_V1_0], [COURIER_V1_1], and [COURIER_V1_2]. A directory named [VERSION] is necessary for the current implementation of the product tools, and contains information about which product is installed. Future implementations will probably centralize this information with the rest of the database.

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



Sample Product Version Directory

The [PRD_V1_0] directory is the Version's Rooted directory, while [COM] and [SYSTEM] are required directories, and [MAINT] is an optional directory. Beyond these directories, the developer can use any tree structure (under the product's version rooted directory).

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 which points to lower level in the directory tree:

```
'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 the first figure, the rooted logical name for COURIER is

```
$ SHOW LOGICAL COURIER$ROOT
"COURIER$ROOT" =
  "disk:[VAXONL.COURIER.COURIER_V1_0.]"
```

Required Files

The product specification requires that each product provide two command files, of defined logical 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.

Developing the Products (BUILD)

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).

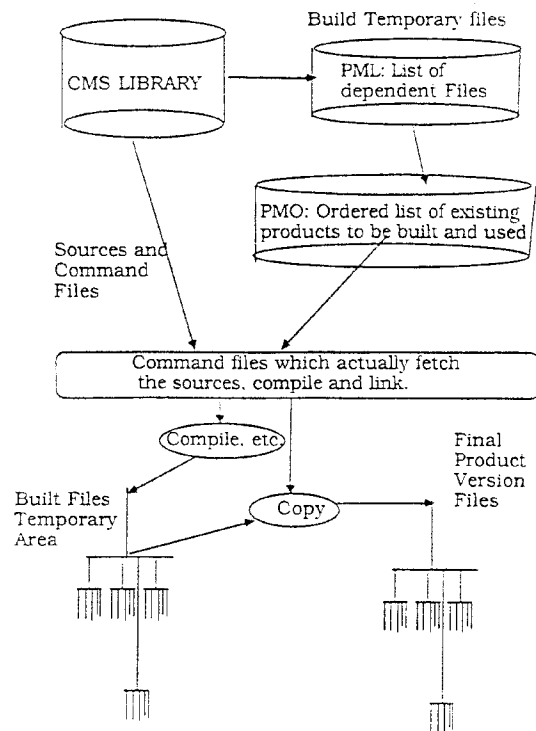
For example, the product KERMIT_VMS is to be "BUILT":

- 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's 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.



Steps in "Build"ing a product

System Management of Products (SITE 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. 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 a Bulletin is posted on the system.

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.

Using the Products (PRODUCT_SETUP)

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.

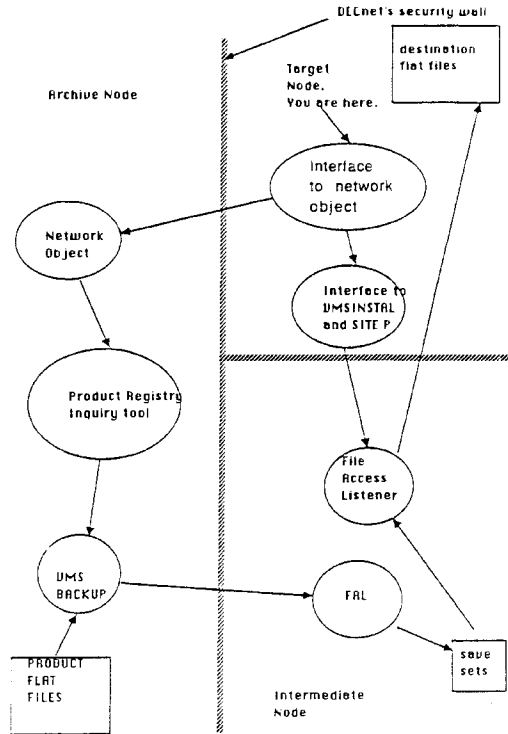
The ability to switch transparently between product versions is provided by the logical name tables created for the product. When switching between product

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.

Obtaining the Products (DISTRIBUTE)

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.



Distribute Processes

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.

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.