

Distributed Computing with Creo Elements/Pro Distributed Pro/BATCH

White Paper

Leverage your hardware and software resources to perform batch jobs over the network

Executive Overview

Distributed computing software has been around for years, yet has remained unutilized by software engineers. Many typical engineering activities, which tend to be time-consuming, compute-intensive, and highly repetitive, benefit greatly from this concept. Companies using Creo Elements/Pro Distributed Pro/BATCH will realize better software license utilization by scheduling and executing batch processes during off hours. This increases users' productivity by allowing them to concentrate computer resources on interactive tasks rather than waiting on non-interactive tasks to complete. Additionally, distributing batch processes across the network to multiple qualified services can reduce overall processing time.

Pro/ENGINEER Distributed Pro/BATCH offers a standard task-based XML framework for creating, configuring, and saving user-defined tasks. Application programming interfaces included with the Distributed Services Manager allow companies to create custom clients using C/C++ or Java, and services using C/C++. These tools allow companies to create custom batch tasks for Pro/ENGINEER Wildfire as well as other software applications.

Creo Elements/Pro Distributed Pro/BATCH Architecture

Customers using Pro/ENGINEER Wildfire or later may leverage the scheduling and batch processing capabilities of Distributed Pro/BATCH. The Creo Elements/Pro Distributed Pro/BATCH architecture shown in Figure 1 is comprised of the following elements:

1. **Distributed Pro/BATCH Client** allows users to create, save, schedule, submit, and monitor batch processing tasks. The ability for Clients to operate within the Windchill environment is available in Pro/ENGINEER Wildfire 3.0 and above.
2. **Distributed Services Manager (DSM)** is a purchasable, platform independent, network-based resource management software system to manage remote resources for performing some service over a network. It includes six Distributed Pro/BATCH Service licenses. An application programming interface is included as part of the DSM in Pro/ENGINEER Wildfire 2.0 and above for creating custom clients and services.
3. **Distributed Pro/BATCH Services** perform the work of scheduled tasks on individual, configured workstations. Each active Service instance requires a Distributed Pro/BATCH Service license. Multiple service instances may exist on a single workstation at one time. Additional Services are purchasable for use with the Distributed Services Manager.




Figure 1: Distributed Pro/BATCH Architecture

Creo Elements/Pro Distributed Pro/BATCH users create and schedule batch processing requests using the Distributed Pro/BATCH Client. The DSM receives client requests as a "task group" in XML format, identifies available network Distributed Pro/BATCH Services capable of executing the requests, and distributes the tasks appropriately. Service instances execute the tasks and return the results to the DSM. The DSM collects and returns the results to the client according to the output location specified in the original XML task request. If a Distributed Pro/BATCH Service fails to execute an individual task, descriptive error codes are written to the task group log managed by the DSM (see Figure 6).

Distributed Computing Technology

Distributed Computing technology has seen tremendous advances by the proliferation of web protocols such as Simple Object Access Protocol (SOAP), XML and HTTP. Creo Elements/Pro Distributed Pro/BATCH and the Distributed Services Manager are built using these latest web standards in keeping with PTC's endeavor to provide customers with quality toolsets built on the latest technologies to optimize their usage of existing resources.

For example, all communication between distributed clients, distributed services, and the DSM are handled using standard SOAP requests/responses over HTTP. All batch-processing instructions are formatted in XML (Figure 2) for ease of use, visualization, and customization.

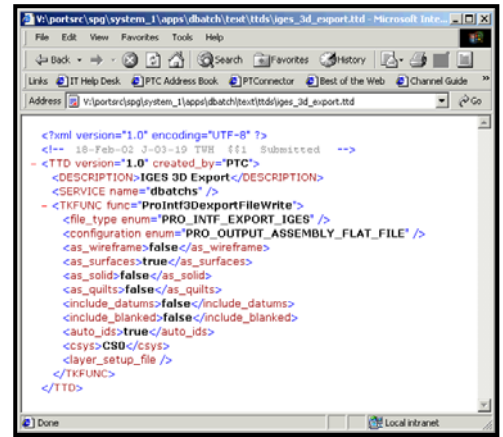


Figure 2: XML-based batch processing instructions

Distributed Pro/BATCH Client

Creo Elements/Pro Distributed Pro/BATCH features a client graphical user interface (See Figure 3) to compose, send and monitor batch jobs. The client can be installed and run independent of Creo Elements/Pro Wildfire. The client also features a scheduling capability.

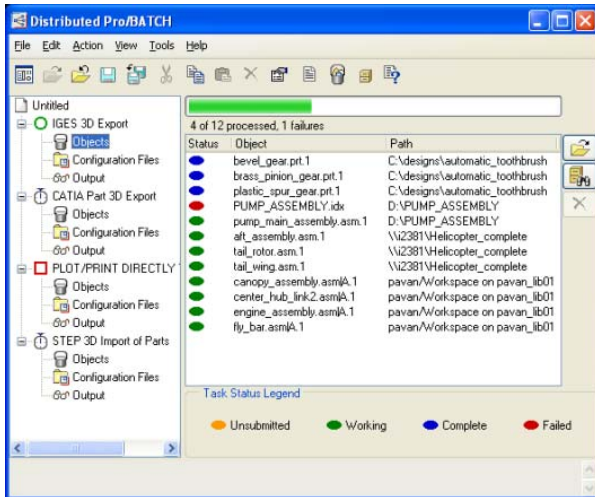


Figure 3: Pro/ENGINEER Wildfire 5.0 Distributed Pro/BATCH Client User Interface

The Distributed Pro/BATCH Client is provided with a series of standard task type definition (TTD) templates in XML format designed to execute out of the box services including batch:

- ATB Update
- Image Generation
- Printing and Plotting
- Execution of ModelCHECK™ analysis
- Pro/ENGINEER Wildfire Model Updates
 - such as save with display
- Exporting and importing from 2D interfaces
 - DXF, DWG, and many others
- Exporting and importing from 3D interfaces
 - STEP, Parasolids, CATIA and many others
- Exporting to visualization formats
 - ProductView, VRML, PDF and many others

In addition to these out of the box capabilities, Pro/ENGINEER Wildfire 4.0 introduced the ability for companies to create Pro/ENGINEER Wildfire custom XML tasks using Pro/TOOLKIT without requiring support from PTC. This capability is included within the DSM and requires a Pro/TOOLKIT license.

Engineers and CAD administrators initiate the Distributed Pro/BATCH Client and load the desired TTD(s). Users provide a list of files to process, optional configuration files like plotter *.pfc and config.pro, and specify the results output location. The fully configured task may be submitted to the DSM. The status icon provides users with visual feedback as objects progress from unsubmitted to working and then to either complete or failed. Tasks may also be saved to XML format (Figure 2) for later processing or future reuse. Saved tasks may be executed from the command line.



Distributed Resource Management with DSM

Distributed Services Manager (DSM) is a platform independent, network-based application introduced in Pro/ENGINEER Wildfire. The DSM allows engineers or CAD administrators to setup and manage a “server farm” of network resources, also referred to as service resources that are used to perform client requested tasks. The graph in Figure 4 shows how a server farm of service resources can dramatically decrease processing time and how performance improves as the size of batch jobs increase. Engineers don’t tie up their hardware resources performing the task.

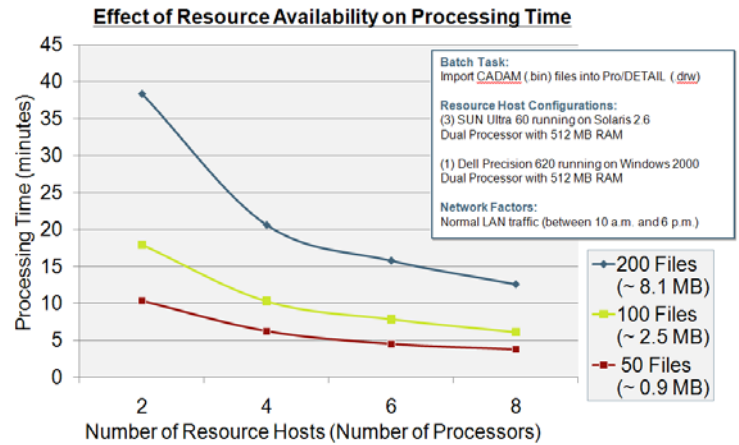


Figure 4: Distributed Processing Scalable Efficiency

CAD administrators may also configure the DSM to create a service bureau. Rather than installing and configuring every Pro/ENGINEER Wildfire client to support a specific process, CAD administrators can configure one commonly available workstation as a Distributed Pro/BATCH Service for all Pro/ENGINEER Wildfire users. This allows companies to better utilize software licensing and facilitates administration. Any client can submit a request to the DSM and that request will be handled by the properly configured and licensed service. If necessary, client requests are queued by the DSM and assigned to available service resources in order.

Companies can leverage this technology to schedule tasks for when service resources will be “available” to process specific requests from remote clients. Users can specify execution of tasks to a later time, allowing more critical batch tasks to be processed in the interim.

CAD administrators may use the DSM to monitor and manage (cancel or re-queue tasks) submitted client requests. The DSM also allows CAD administrators to monitor and manage (specify, schedule, enable or disable, start or stop) Distributed Pro/BATCH Service resources. The DSM maintains a log of all activity and reports descriptive error codes for many common task failures.

A web-portal interface to the DSM (Figure 5) provides users and CAD administrators with online access to monitor the progress of their jobs from any location. Creo Elements/Pro Distributed Pro/BATCH is the first distributed service developed by PTC to leverage this distributed resource management capability.

Task ID:	User Name:	Client Host:	Service Type:	Service Instance:	Task Status:
1	mlobo	MFLOBO03L.ptcnet.ptc.com	dbatches	1184@mlobo03l	Complete
2	mlobo	MFLOBO03L.ptcnet.ptc.com	dbatches	13800@quiet	Working
3	mlobo	MFLOBO03L.ptcnet.ptc.com	dbatches	13801@quiet	Working
4	mlobo	MFLOBO03L.ptcnet.ptc.com	dbatches	1184@mlobo03l	Complete
5	mlobo	MFLOBO03L.ptcnet.ptc.com	dbatches	1184@mlobo03l	Working
6	mlobo	MFLOBO03L.ptcnet.ptc.com	dbatches	None	Pending
7	mlobo	MFLOBO03L.ptcnet.ptc.com	dbatches	None	Pending
8	mlobo	MFLOBO03L.ptcnet.ptc.com	dbatches	None	Pending

Figure 5: DSM Web Portal

Distributed Pro/BATCH Services

Services are the workhorses of the server farm. Each service is designed to perform a unique task like printing/plotting, or export to ProductView. One workstation can be configured to execute multiple services and many workstations can be configured to execute the same service, offering a very flexible and scalable architecture.

Client requests identify the type of service required and the DSM monitors services for availability and submits a processing request to an appropriate, available service. A Distributed Pro/BATCH Service license is consumed when an instance of the Service is initiated to process the request. When the task is complete, the service instance returns the results to the DSM and closes. This means a company with six service licenses may install more than six Distributed Pro/BATCH Services. However, only six service instances may be executed simultaneously.

Distributed Pro/BATCH Logging

Confirmation of success, notification of warnings, and explanation of failures are invaluable to users and administrators. Distributed Pro/BATCH Clients receive notification of successful and failed tasks directly in the client user interface (See Figure 3).

Administrators can review log files for the DSM and the Distributed Pro/BATCH Services (see Figure 6). Distributed Pro/BATCH Services maintain task instance logs, one for each task instance processed. For performance reasons, task instance logs must be enabled and they do not persist after a task instance. The DSM maintains a task group log for each client submitted task group. The task group log records task instance successes and failures. Pro/ENGINEER Wildfire 4.0 (M050) introduced more granular failure reporting. See the About Logging in Distributed Pro/BATCH documentation for additional details.




Figure 6: Log File Locations

Creo Elements/Pro Distributed Pro/BATCH Operating Modes

There are two modes, or implementations of Creo Elements/Pro Distributed Pro/BATCH: Standalone and Distributed.

1. **Standalone:** In standalone mode, a local session of Pro/ENGINEER Wildfire on the user's machine is used to process batch jobs. Scheduling is extremely useful for the standalone mode, since the user can schedule when the local session of Pro/ENGINEER Wildfire will start processing batch jobs. Standalone mode is included in the Pro/ENGINEER Wildfire base package, is selected for installation by default and does not require licensing or installation of the DSM or Services.
2. **Distributed:** In distributed mode, one or more remote sessions of Pro/ENGINEER Wildfire on the network will process the batch jobs. One or more distributed Pro/BATCH clients on the network may submit batch jobs to be processed. A centralized DSM manages all batch requests and services. The distributed configuration requires installation of the DSM, configuration, and a requires licensing of the DSM and Services.



Creo Elements/Pro Distributed Pro/BATCH in Use

Below are a few examples of how Distributed Pro/BATCH can be applied in order to make use of available resources to perform trivial tasks that are otherwise laborious.

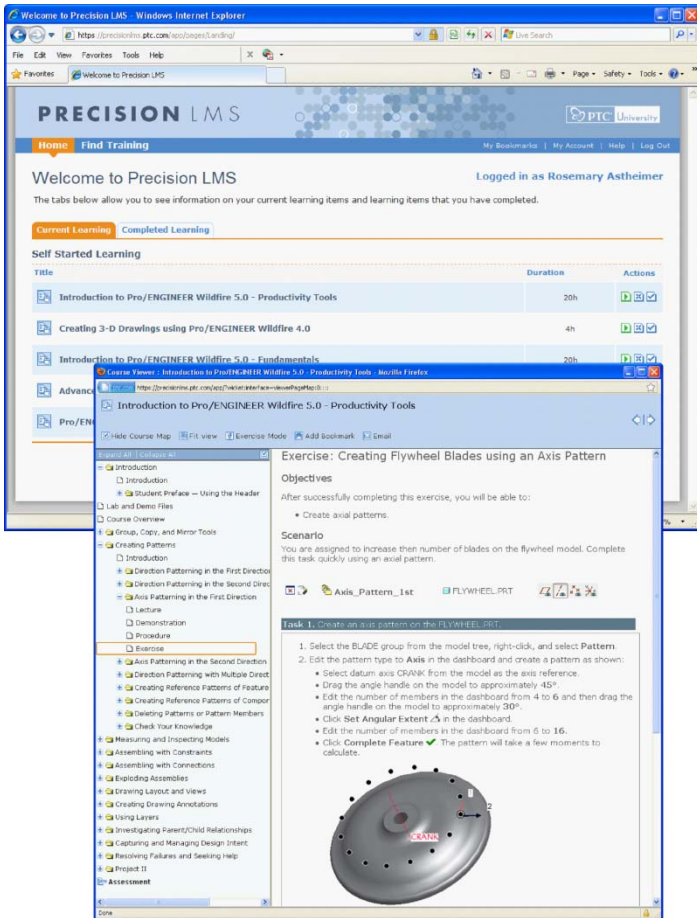


Figure 7: PTC Precision LMS

PTC Precision Learning Management System Grading Engine

PTC University has integrated Distributed Pro/BATCH into their learning management system to grade Pro/ENGINEER customer models. This helps PTC to verify the effectiveness of its training materials and the users to assess their understanding of the topics learned.

European Auto Manufacturer

Batch Translation of 1,500 - 2,000 Design Models between CATIA and Pro/ENGINEER each night running Distributed Pro/BATCH "Out of the Box".

Heavy Equipment Manufacturer

Proposed translation of Pro/ENGINEER data from approximately 200 Pro/ENGINEER users per site at 16 sites worldwide to enterprise visualization and document formats Data using Distributed Pro/BATCH "Out of the Box".

Medical Product Manufacturer

Custom TTD used to perform a batch addition and/or removal of parameter data from Pro/ENGINEER models.

Asian Auto Manufacturer

Batch testing of 60,000 legacy Pro/ENGINEER drawings to check for user modeling errors via a custom TTD.

US Truck Manufacturer

Batch upgrade of 120,000 Pro/ENGINEER library part models to an updated version using a Custom TTD.

Office Equipment Manufacturer

Batch conversion of product models to web-friendly format for customer-access catalog. The use of DS APIs to perform this task increased throughput by 300% while saving \$0.5M per year.

Summary

Creo Elements/Pro Distributed Pro/BATCH offers robust, scalable capabilities in the area of batch processing of common Pro/ENGINEER Wildfire related tasks, such as executing ModelCHECK, plotting or exporting ProductView viewables. Distributed computing technology in Pro/ENGINEER Wildfire leverages the latest open web standards, namely SOAP, XML, and HTTP. Creo Elements/Pro Distributed Pro/BATCH in conjunction with the Distributed Services Manager for network-based distributed resources management heralds a new paradigm of batch processing with Pro/ENGINEER Wildfire.