

Teamcenter multi-site collaboration

White Paper

Enabling integration of global design and manufacturing sites and the supply chain.

www.siemens.com/plm

Contents

Executive summary	3
Teamcenter multi-site collaboration module	5
Fault tolerance.....	6
Scalable, flexible multiple-site deployment.....	6
Multiple model capability	6
Optimal leveraging of networks and servers	7
Power of multi-site collaboration linking...enabling collaboration.....	7
Advantages of replicating versus exporting	7
Supplier integration.....	8
Characteristics of a successful implementation	10
References	10

Executive summary

Today's increased focus on revenue growth and international competitiveness has driven manufacturing companies to seek new ways to accelerate product development through process innovation, global alliances and strategic partnerships. Shifting business and market dynamics demand an agility undreamed of – even a decade ago. But along with the new opportunities, globalization has also created tough new challenges.

A recent report published by Deloitte analysts highlights some of the difficulties manufacturers need to overcome to operate globally: “Our study found that the clear majority of manufacturers’ value chains lack the capabilities to effectively generate new products and sustain them over their lifecycles. One major reason is the rapidly increasing complexities of global markets and value chains. To reduce costs and pursue global markets, companies are dispersing engineering, design and manufacturing, as well as marketing, sales and service activities around the world. This has made it significantly more difficult to coordinate the value chain and manage products effectively across their lifecycles.”

In a fast-paced, globally-distributed environment, all team members need immediate access to all types of product information. At every stage of a product’s lifecycle, they need to be confident that the data they access is both accurate and up-to-date. An example of why this is so important can be seen in the need for each team member in the extended enterprise to be able to see design changes and their impact in a specific design context; as well as view their impact on the overall product in the context of the program.

If suppliers deliver products based on outdated information, costs go up and time-to-market increases.

Each location in a global company may want to optimize its product database locally, which limits the database’s relevance to the extended enterprise. Often, this lack of relevancy occurs because the database can only be accessed locally. However, if the information is organized properly, as well as managed and distributed to all who need access, it will provide material for effective decision making and help the organization achieve its next level of competitive advantage.

Consider the example of FEI, a mid-sized company based in Holland, which makes high-tech laser machines. “The specialties of our different global sites complement each other, but working together often resulted in confusion and therefore inefficiencies,” explains Frank de Jong, research and technology manager at FEI. “We realized that if we were

better integrated, we could profit more from the varied product portfolio as well as our mutual architecture and innovations. That would also reduce time-to-market.”

Geographically dispersed teams need easy access to workflow-driven process information and the ability to re-use information from earlier delivery programs – because those abilities make it possible for far-flung enterprises to reduce costs and resource consumption. Along with facilitating component re-use, manufacturers also require common access to design processes to accommodate evolving product configurations along with their options and variants.

Take the example of Bosch-Siemens Household Goods (BSH), the Munich, Germany-based manufacturer of world-class consumer appliances. The company competes in global markets and addresses local opportunities by establishing design and manufacturing sites around the world. “The purpose of a PLM system is that everyone should use it in the same way. If you are designing a parts family, an engineer needs to number the parts in a standard way so another engineer in another facility can recognize the parts for use in a product,” says Uwe Tontsch, head of digital engineering at BSH. “If a product is designed in Germany, but needs to be modified in North America or China for local markets, the engineer needs to know what he is looking at.”¹

Geographically dispersed teams need tools for facilitating teamwork, and require new systems and processes to transcend geographic and organizational boundaries.

Establishing a managed global product development structure provides manufacturers with the ability to make decisions based on real, true and current information. It helps the company, working teams and individuals spend less time on nonvalue-added activities, such as finding the latest version of the product or hunting for the right design context. It enables manufacturers to reduce risk and time by intelligent component re-use and makes it possible to validate product changes faster, more reliably and earlier. Most important of all, it provides the agility necessary for the enterprise to address new market opportunities quickly.

The implementation of product data management (PDM) systems throughout the extended global enterprise has been notoriously difficult in the past because of the conflict between providing access to widely scattered users and maintaining a single copy of each data object. Various approaches have been taken to try to address this conflict, each trying to tackle portions of the problem.

It's no longer uncommon for the design tasks on a single product to be dispersed to engineering groups located on different continents. In order for these groups to work effectively on the same project, engineering data needs to be distributed throughout the organization – to all engineering sites, manufacturing sites and typically to external suppliers – with the integrity of each facility's changes permanently maintained.

The biggest challenge in achieving global collaborative development involves distributing complex product structures and allowing complex changes while still maintaining data integrity. When looking at the various activities that need to be supported in a collaborative environment, the minimum requirement of concurrent design is that users throughout an organization at least be able to view objects that are relevant to their own work. To effectively support collaborative design, users at different sites should not only be able to view the objects, but also the relationships between those objects and others throughout the extended enterprise. The system must be able to understand the rules associated with the relationships and support the behavior required by those relationships. The system must aid the user in building products that are valid with regard to their usage throughout the organization.

In summary, the critical requirements for a global product development network are:

1. Performance – The solution must satisfy minimum performance criteria
2. Support of different network topology, including local-area network (LAN), wide-area network (WAN) and mobile – Some users may be on LAN, some remote sites may be on WAN and some users may be only occasionally connected to the network on a mobile device. The solution must be able to provide the right user interface, performance characteristics and connectivity to support all types of users
3. Security – Enterprise security must be enforced at all locations
4. Data integrity – The system must ensure that only the master instance of an object is allowed to be modified
5. Fault tolerant – An information technology (IT) outage at one location (including the headquarters data center) should not impair other locations from functioning
6. Scalability – The enterprise should be able to grow the network incrementally and to support large numbers of users
7. Flexibility – Enterprises must be able to add and remove a site easily

Teamcenter multi-site collaboration module

To implement a PDM system that can be used to successfully manage product content across multiple sites, there needs to be a high degree of deployment flexibility to implement a solution in phases so that each site can deploy when ready. It must be independently deployable at each site, yet be able to function throughout the global enterprise as these divisions start to work together.

Although there are numerous requirements for a global enterprise-wide solution, two major requirements should be addressed. The system must satisfy user-level and administrative-level requirements.

User-level requirements determine how users view and access the system, as well as reflect specific needs such as the ability to locate information anywhere on the network, view data from multiple sites simultaneously and provide integrity (even when manipulating data at multiple sites at

the same time). The system should be able to be used to create intelligent references between objects in multiple systems, build product structure using assemblies and subassemblies from multiple sites and guarantee that users always have up-to-date data.

Administrative-level requirements include the ability to upgrade databases independently, support local site autonomy and facilitate comprehensive participation in the enterprise solution. The system must provide excellent and consistent performance for all users in the network, as well as data security for both networked and individual sites. This includes security for individual objects on the network.

When a company decides to implement a PDM solution, one option it normally considers is to allow individual divisions to implement the system while gaining a better understanding of how the system works. A second option is to define the

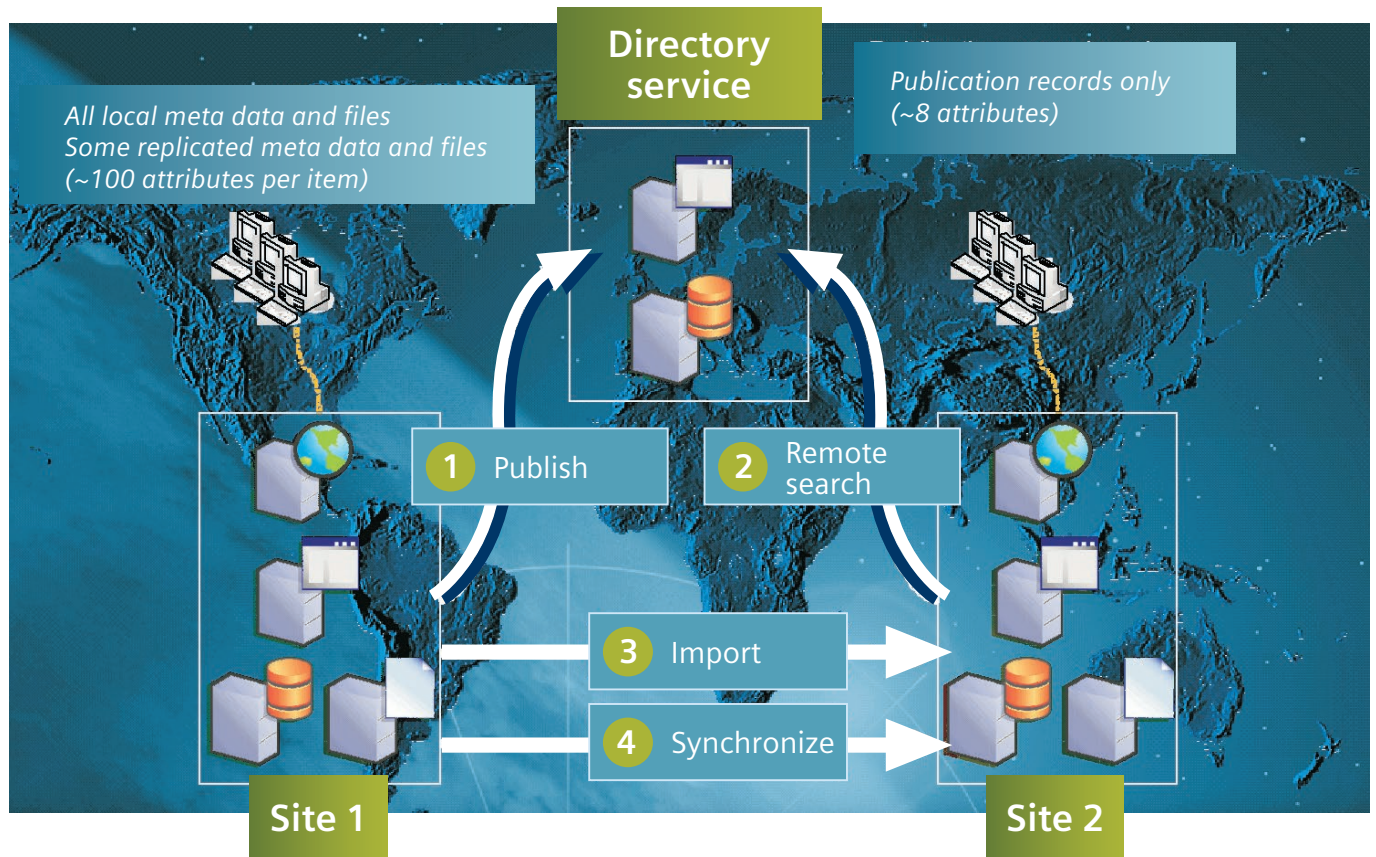


Figure 1: Global multi-site fundamentals.

architecture for the entire corporation and build the complete corporate system as a whole. The former option allows companies the luxury of developing a basic plan and laying the groundwork for a large-scale deployment with minimal risk, while the latter requires a great deal of planning and coordination throughout the corporation before beginning deployment. Even with a great deal of upfront planning, situations, divisions, networks and costs change and international considerations often become an issue.

Teamcenter® multi-site collaboration software allows each site to have its own resources, models, sets of users, groups and business rules. Business rules consist of access controls, release procedures, change management practices and other extensions to the system. When running in a global Teamcenter environment, the systems can be managed and maintained independently or cooperatively. This means that with Teamcenter local configurations and customizations are possible, and different versions of the software can be deployed yet still interact with other versions of the software.

Fault tolerance

One major advantage of this architecture is that it is fault tolerant.

That means that if data cannot be retrieved due to short-term network problems, the system can gracefully recover and allow the user to access the information later. The architecture provides for continuous operation. If a particular node or database is not reachable on the network at a given time, users can continue working on local copies. Only the small number of users who are completely dependent on a particular database would be without access in this situation.

Scalable, flexible multiple-site deployment

The Teamcenter architecture allows individual divisions to implement autonomously while providing the groundwork for a large-scale implementation. In fact, individual sites can be implemented with no knowledge of one another. As the sites decide to coordinate their activities, they simply set up a directory services database and notify one another of their existence. This avoids extensive early planning while minimizing the risk involved in the big bang approach required by alternate architectures, in which a firm is forced to define the entire implementation upfront.

Teamcenter multi-site collaboration allows objects to be shared among multiple sites. However, it only shares objects that are specifically requested, thereby maximizing the use of company resources by minimizing the need to replicate huge numbers of objects that may never be used at those sites.

By using an intelligent, controlled replication scheme, Teamcenter multi-site collaboration allows enterprise business rules to define what information should be shared, and by whom and when. It makes needed information easily

available without compromising the security of the participating sites, which is especially important when working in a customer/ supplier environment.

By using an object-based replication scheme, Teamcenter multi-site collaboration not only minimizes use of the network by users accessing information from various sites, but also optimizes the performance for users doing their normal day-to-day activities.

Companies can implement individual sites autonomously. If two or more sites decide to coordinate activities with one another, they simply set up a directory services database for the participating sites and begin sharing information. Each site continues to maintain its own business rules, its own set of users and groups; alternatively, the site can adopt rules and business objects from the other sites. An access control capability (for the various objects created and managed within a site) manages which objects are permitted to be viewed or modified outside the site. Rules are controlled by individual sites. When objects are shared with other sites, the object keeps track of which sites have (read-only) copies of the information and provides those sites with updates if and when necessary. This high-level replication scheme allows objects to be controlled at the semantic level rather than at the raw data level, thereby providing more control over complex objects such as entire product structures. In addition, more sites can be easily added.

Multiple model capability

A product data management implementation starts with the data the system must manage. The global Teamcenter system is defined by both a set of protocols through which the participating systems share information and the model from which they communicate. As with other systems, the protocols must be well defined, and the models must have a certain commonality that all systems understand. But with Teamcenter, the physical models may differ, allowing local customizations at individual sites, as long as the systems have an understanding of what they are managing and what it means to have relationships among individual objects.

As stated earlier, the primary role of a PDM system is to manage product content and guarantee the integrity of that information throughout the extended enterprise. That's why when moving to a distributed implementation containing information from multiple systems, the Teamcenter multi-site collaboration system is so effective. It is specifically designed – and enabled by its architecture – to manage more than simple references among systems in independent databases. It guarantees the integrity of the information throughout those databases, enforcing not only the physical relationships between those objects, but also the semantic, or business, relationships between those objects. With Teamcenter multi-site collaboration, products being built from assemblies created and managed at multiple sites fit together when completed. Changes that affect those

assemblies in the context of the main assembly are shared with all affected assemblies in a timely manner, allowing designers to evaluate such changes and take corrective action before problems reach the production line.

Optimal leveraging of networks and servers

When implementing PDM on a global scale, it is important to consider the deployment of networks and computers – both indispensable, high-impact resources.

Networks can represent a make-it or break-it factor when establishing multiple systems throughout an enterprise. Networks vary with respect to their speed and reliability. On any given day, the route that information takes through a network can be different; giving random performance characteristics for any product that is network-intensive. If a network segment goes down, it is imperative that such failures do not impact the ability of users to do their jobs. If it does affect users, it should impact the fewest number of people possible.

Server availability is another important consideration. Servers go down for various reasons. Use of certain computers can be so intense that they are effectively unavailable as well. Backup schedules may require that computers be taken offline for extended periods of time. It is important to ensure that these types of problems do not affect the overall productivity of an organization. Multi-site collaboration maximizes the use of a company's existing network infrastructure when replicating and synchronizing data across sites. Unlike traditional file movement technology based on the use of Transmission Control Protocol (TCP)/Internet Protocol (IP) or File Transfer Protocol (FTP) that demonstrates an exponential decay in performance as network latency increases, Siemens PLM Software uses its file management systems (FMS) software to move metadata as well as computer-aided design (CAD) designs, Microsoft Office documents and other file-based information between sites.

FMS delivers superior performance in comparison to moving information via FTP.

FMS utilizes multi-channel/multi-source streaming technologies to break a file up into multiple parts, transmit them in parallel and then reassemble them at the other end. This reduces the time taken to transmit a file. The higher the network latency and the larger the file, the greater the gains when compared with the FTP alternative. Typically, gains are two times at 100ms through to four times at 500ms.

- FMS supports volume-to-volume information streaming between sites rather than having to move a file using multiple hops as is typically the case when information is moved from one managed environment to another. In a non-FMS environment, you would typically see the following process: export file from the locally managed area to a local working area, copy the file from the local

working area to a remote working area and import the file from the remote working area to the remotely-managed environment

- FMS exhibits linear performance decay as packet loss and latency increases, as opposed to the exponential performance decay that occurs with FTP

Teamcenter multi-site collaboration minimizes the use of networks in sharing information between multiple sites. As earlier stated, with the concept of intelligent replication, Teamcenter enables you to optimize the use of networks by keeping most of the work done by users at the local database and only going to the network to replicate objects. If a network or system goes down, or computers with network access become unavailable, users can continue to work on information that is local, then resume working on the remote objects when the network or remote system is brought back online.

Some solutions are designed to access the network every time they need an object from a remote database. This increases network traffic and requires that both the network and the remote system be up and accessible at any given point in time. Such dependencies often cause system performance problems or system hang ups if one of the networks is down.

Power of multi-site collaboration linking...enabling collaboration

The Teamcenter architecture facilitates collaboration at multiple levels by allowing subassemblies to be worked on while maintaining links between those structures at multiple sites. The engineer who is the owner of the objects is able to create business rules that define how those subassemblies or components interface together and are manufactured. These rules are shared and enforced throughout the enterprise, guaranteeing that changes made at one site are valid in the context of work being done by other parts of the organization. Compatibility of shared assembly changes is propagated to all parent assemblies, and change control is coordinated among all users of those parts.

Advantages of replicating versus exporting

Some systems export a copy of the object for local modifications and reimport the new version back to the original database to save it, making the edits outside the context of the PDM system. The best systems allow replication of data to local sites. However, the use of replication raises data integrity, security, auditing and tracking issues that must be addressed. Once an object is exported from the PDM system, it is free to be shipped anywhere in the world. Without proper security controls, replicated product information could fall into the hands of unauthorized people. A replication-based system, which is preferable for its many advantages, must also track replicas of objects in order to ensure that they are safe and secure no matter where they are sent.

Supplier integration

Teamcenter multi-site collaboration can be deployed in a variety of ways to support suppliers. A supplier site could be hooked up just like any other site with direct access into the network in a peer-to-peer configuration. However, for more secure access, Teamcenter provides two other configurations – hub and proxy.

A hub configuration enables secure and seamless integration of partners and suppliers by employing a data-sharing hub as a distribution center.

The hub site is designated as a distribution center for sharing data among cooperating sites. All shared objects are replicated at the hub and automatically published to the hub’s object directory service (ODS). Any user in a cooperating site who wishes to get a copy of a shared object would search the hub’s ODS and then get a copy from the hub.

The hub configuration isolates the internal sites from the external sites, thereby enhancing network security. It also eliminates the need to directly connect external sites to the individual internal sites.

The proxy server configuration is a special hardware/ software configuration that enables secure and seamless integration of suppliers. Typically, a proxy host located inside a corporation’s firewall will be installed with the Teamcenter proxy server software. All suppliers outside the firewall communicate only with the proxy host; there is no network connection to the internal hosts.

The proxy server acts as a relay between the internal sites and the external sites. This basically isolates the internal sites from the external sites, thereby enhancing network security. It also eliminates the need to directly connect external sites to the individual internal sites.

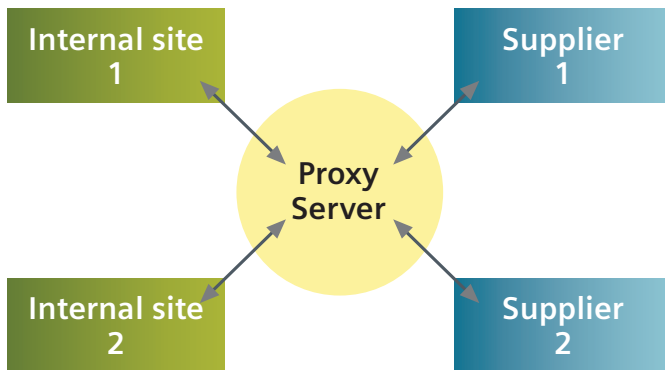


Figure 2: Proxy server.

In an original equipment manufacturer(OEM)/supplier environment, it is not always desirable to have suppliers tie into all the sites inside the firewall of an OEM. In order to provide more restricted access through a firewall, the use of a hub site is recommended. A Teamcenter hub is a site that is

designated as being able to share copies of copies with other sites. This allows users to send a copy of an object to the hub, and allows suppliers to take copies outside the organization. The hub keeps track of where it sends copies, just like the owning site. To the supplier, it appears that the hub is actually the true owning site of the object. The rest of the network is hidden from the supplier.

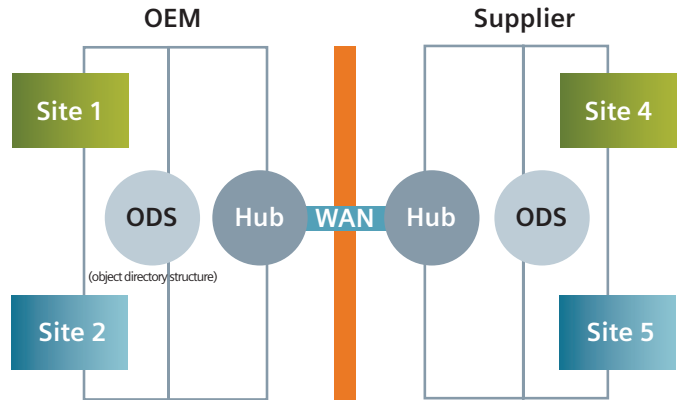


Figure 3: Supplier integration through a hub.

In this configuration, the hub is the only connection to the outside world for the supplier and provides a single opening in the firewall for the supplier site. In the above figure, site 4 can act as a hub for the supplier and allow the supplier to implement its own multi-site collaboration deployment and participate in a multi-site collaboration deployment with the OEM. Figure 3 illustrates this concept.

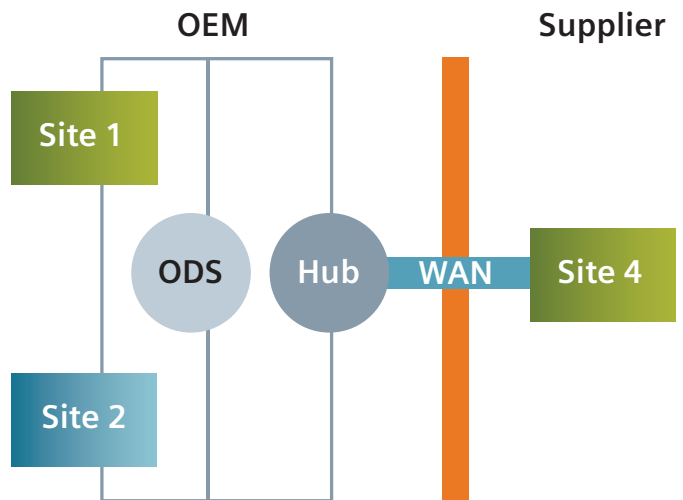


Figure 4: Supplier multi-site collaboration deployment.

In the configuration in figure 4, the supplier has a full Teamcenter multi-site collaboration deployment and creates its own objects, but is able to bring in objects from an OEM and share them in its extended enterprise consisting of multiple sites. This allows the supplier to build components in the context of the primary product, and to share its components with the OEM so that the two companies can work very tightly together, update each other and help guarantee the accuracy of their shared product.

This environment can be extended with multiple suppliers, and the supplier can actually act like an OEM and build its own distributed supply chain using the same technology. Thus the environment depicted in figure 5 can be multilevel, with a larger and larger extended supply chain.




Figure 5: Teamcenter multi-site collaboration in an extended supply chain configuration.

Characteristics of a successful implementation

Building a distributed PDM environment involves many considerations and requirements. Foremost is the power of the architecture in delivering data integrity and maintaining the semantics of the information between the product's various pieces, even if the product is stored in multiple physical repositories. The second consideration is the ability to add new sites as required and enable each site to change and grow over time.

Reliability, availability and consistent performance are crucial to successful deployment, with minimization of corporate resources (such as network bandwidth) also being important. System maintainability is often overlooked.

Moreover, in most organizations, the ability to deploy multiple versions of the software and have them interoperate is imperative. Easy access to information, such as availability through the web and on the desktop, greatly improves user acceptance.

The use of Teamcenter demonstrates the key attributes of successful implementation through architecture, process-based technical breakthroughs and real-world success. Teamcenter is currently in production at multiple sites globally, and is based on technology that has been proven in large-scale implementations over several years.

References

1. "Design Globally, Act Locally" *Managing Automation*, June 25, 2005

Siemens PLM Software

Headquarters

Granite Park One
5800 Granite Parkway
Suite 600
Plano, TX 75024
USA
+1 972 987 3000

Americas

Granite Park One
5800 Granite Parkway
Suite 600
Plano, TX 75024
USA
+1 314 264 8499

Europe

Stephenson House
Sir William Siemens Square
Frimley, Camberley
Surrey, GU16 8QD
+44 (0) 1276 413200

Asia-Pacific

Suites 4301-4302, 43/F
AIA Kowloon Tower, Landmark East
100 How Ming Street
Kwun Tong, Kowloon
Hong Kong
+852 2230 3308

About Siemens PLM Software

Siemens PLM Software, a business unit of the Siemens Digital Factory Division, is a leading global provider of product lifecycle management (PLM) and manufacturing operations management (MOM) software, systems and services with over nine million licensed seats and more than 77,000 customers worldwide. Headquartered in Plano, Texas, Siemens PLM Software works collaboratively with its customers to provide industry software solutions that help companies everywhere achieve a sustainable competitive advantage by making real the innovations that matter. For more information on Siemens PLM Software products and services, visit www.siemens.com/plm.

www.siemens.com/plm

© 2015 Siemens Product Lifecycle Management Software Inc. Siemens and the Siemens logo are registered trademarks of Siemens AG. D-Cubed, Femap, Fibersim, Geolus, GO PLM, I-deas, JT, NX, Parasolid, Solid Edge, Syncrofit, Teamcenter and Tecnomatix are trademarks or registered trademarks of Siemens Product Lifecycle Management Software Inc. or its subsidiaries in the United States and in other countries. All other logos, trademarks, registered trademarks or service marks belong to their respective holders.

6369-Y8 8/15 C