BIOGRAPHICAL INFORMATION

David J. Peters
Manager, Systems Integration Department
Environmental Systems Research Institute, Inc.

Specific Responsibilities

Mr. Peters is responsible for development of ESRI Systems Integration services, system design methodology and system sizing models, ESRI software performance validation testing, and for development of infrastructure architecture standards for supporting GIS enterprise operations. He also coordinates system architecture design training programs, supports hardware vendor product validation and development of joint hardware technology strategies, and oversees system architecture design and system performance testing services.

Mr. Peters is author and publisher of the System Design Strategies white paper, maintaining a standardized methodology for design and deployment of enterprise GIS operations for ESRI worldwide customers. He also developed and maintains the System Architecture Design Strategies for GIS training class, which is part of the ESRI education curriculum. In addition, Mr. Peters provides key presentations on system design and implementation strategy in support of user conference workshops and pre-conference seminars worldwide.

Past Experience

Mr. Peters has more than 19 years of experience in systems integration management including technical design and system testing. In 1990, he joined Environmental Systems Research Institute, Inc. (ESRI) as manager of the Systems Integration Department for geographic information systems (GIS) after twenty years of experience with the US Air Force.

During his Air Force career, Mr. Peters served four years as chief of the Small ICBM System Test Division assigned to the Ballistic Missile Organization at Norton AFB, CA, where he was responsible for missile processing and launch pad operations for the first missile flight test, ground testing, and nuclear survivability testing of Air Force Small ICBM weapon systems. Mr. Peters also spent 5 years as Air Force lead project officer for the W80 Nuclear Warhead Project Officers Group when assigned to the Air Force Weapons Laboratory at Kirtland AFB NM, where he was responsible for integration of the Air Force Air Launched Cruise Missile and Navy Tomahawk Cruise Missile programs with the Department of Energy W80 nuclear warhead. Mr. Peters also spent 12 years as an Air Force Pilot, including 8 years as a C-141 transport pilot and 2 years as a T-38 Instructor Pilot.

Mr. Peters earned his Master of Science degree in physics from the Air Force Institute of Technology in 1978 and his master's degree in business administration from Southern Illinois University in 1979. He received his Bachelor of Science degree in math and physics from Wisconsin State University–Whitewater in 1970.

PLANNING FOR PRODUCTIVE ENTERPRISE GIS OPERATIONS

David J. Peters

Environmental Systems Research Institute, INC

380 New York St. Redlands, CA 92373 USA

ABSTRACT

Productive GIS Enterprise Operations don't just happen. Successful operations result from proper planning and effective implementation strategies. System performance tuning begins during the planning phase and continues throughout production. An early understanding of the technology, user requirements, GIS information product needs, and software limitations can save money. Planning on a periodic basis (operational planning updates) supports optimum evolution of effective enterprise operations.

More GIS technology is available today than ever before. Network infrastructure limitations introduce the greatest challenges to current GIS enterprise deployments. System architects must understand the limitations of the existing network infrastructure. System administrators must recognize these limitations and plan now for future technology infrastructure demands. This paper provides some guidelines to support customers in addressing growing network infrastructure needs.

GIS is a compute intensive software environment. Selecting the proper hardware is critical to deploying a productive operations environment. Maintaining a productive environment and keeping up with technology change is critical in supporting effective operations.

Appropriate system testing and tuning supports the health and performance of enterprise operations. Tips on when and what to test, along with a close look at critical system performance factors, provide forward looking guidance on how to establish and maintain productive system performance. Understanding the technology and proper periodic planning can produce effective enterprise GIS operations.

PLANNING FOR SUCCESS

Understanding existing business processes and where GIS will best support user workflows is an important first step in developing a productive Enterprise GIS environment. Data requirements and application needs should be derived from understanding the information products and workflows required to support normal business processes. The GIS implementation schedule is often determined by the time required to develop or collect required GIS data resources and to deploy appropriate GIS applications to support user workflows. Estimated peak user loads are established based on the operations deployment schedules (when the system is ready to support user operations and the users are deployed).

Establishing hardware requirements based on user workflow needs is fundamental to supporting productive operations. Technology is changing very rapidly, and proper investment in the hardware infrastructure is critical in supporting user performance needs. There are many ways to deploy a GIS workflow solution, and many of those ways lead to ineffective operations. It is important to understand what technology will support operations within available infrastructure limitations, and build a system solution that will perform.

System architecture design can be viewed as a performance chain. This performance chain is made up of several components including user workflow, database design, platform selection, network capacity, and data storage. User productivity is determined by the weakest link in this system design. The weak link can establish a performance bottleneck that must be resolved to support user performance needs. System environments should be designed so current CPU processing technology is the limiting design component.

Proven methods have been established over the years to support development of proper deployment strategies and architecture design selection. A clear understanding of user workflow requirements provides a foundation for the system architecture design process. The system design process should include a review of current technology, the existing customer environment, and system configuration alternatives. Peak user processing loads generate requirements on the selected system components, and these system component loads are translated into appropriate hardware specifications. Hardware should be purchased just in time to support operational deployment needs.

UNDERSTAND THE TECHNOLOGY

Selecting the right software technology to support user workflow requirements is a very important decision, and can have significant user performance implications. GIS technology has evolved over the past 20 years to support a broad variety of user workflow requirements. Technology decisions early in the development specify the selected GIS data source (Database, File Systems, Web Services) and how the application is deployed (ArcGIS Desktop, ArcEngine, Web Applications and/or services, and mobile disconnected environments). ArcGIS Desktop may be distributed on the user workstation or deployed on terminal servers supported by central Window Terminal Servers. Web applications and services may be supported by ArcGIS Server or ArcIMS technology. The technology selection determines the application functions available to support the user workflow; technology selection also determines platform and infrastructure loads generated to support user workflow operations.

IT organizations are finding that they can effectively support most GIS operations much more effectively from a centralized server and database architecture. Centralized operations support the data and application servers in the central computer facility, significantly reducing overall hardware and administration costs, and reducing implementation risk; significantly improving the ability to support integrated business operations, improved data access, and enhanced system security. Centralized application environments also significantly reduce overall network traffic loads by keeping a majority of the high network

traffic in the computer room and limiting user interface to a much lower bandwidth display environment.

At the same time, many organizations find a need to support distributed GIS operations. Primary reasons for distributed operations include organizational preferences, management confidence, department level physical security issues, and potential infrastructure limitations. Distributed operations cost more and introduce a variety of implementation risks that need to be taken into considerations to support a successful implementation.

Selecting the optimum software technology and appropriate supporting architecture solution will establish a foundation for productive GIS operations.

FOLLOW TECHNOLOGY TRENDS

Much can be learned from the success of others. There are many things we learn about GIS technology by watching what successful customers do.

The evolution of GIS technology provides some clear messages that can help direct us toward the most appropriate system architecture strategy. We learn that GIS resources are more valuable as they are shared with others, and more data resources improve the value of our GIS products. We also learn that consolidated operations are more cost effective and easier to support than those supported by distributed architectures. We also know that the best data comes from the source, so it is important to distribute data custodian responsibilities to the appropriate organizational level. Also, we have learned that our information product outputs are only as good as our data inputs – so quality and currency of our data are important.

History is a good teacher, and we need to listen to the message. Technology is changing very rapidly – Web services, Federated Systems, and Server Oriented Architecture are taking us to a new level. Technology is rapidly approaching a point where we may be able to have our cake and eat it too (support centralized data and applications maintained by distributed users, all with the same optimized infrastructure strategy and secure operations). Soon many of us can go to work without leaving home, meet with customers throughout the country without leaving the neighborhood, and travel throughout the world without leaving town. What used to be a Buck Rogers dream is now reality. Technology is taking us to a new level.

CONSERVE NETWORK RESOURCES

The things we can do with technology are rapidly exceeding the infrastructure needed to do it. We all live in a world of shared resources, both natural and man made. This applies not only to the air we breathe and the water we drink; it also applies to the highways we drive on and – in our new virtual workspace - shared Internet and Intranet communication networks.

GIS is a powerful software technology that can consume large quantities of data and rapidly translate these data into exciting user friendly information products. The most common GIS data sources are shared Geodatabase repositories, accessed directly by GIS users throughout

the organization. As a result, standard GIS operations generate large volumes of data traffic over shared network environments. Understanding the GIS architecture alternatives and the associated infrastructure cost provides critical information needed to support the proper system architecture solution. Complying with recommended design guidelines can have a big impact on user performance; often the consequence of not considering existing infrastructure limitations when developing the conceptual architecture design.

PROTECT YOUR DATA RESOURCES

Data are the most valuable resources supporting effective business operations. Loss of data can have drastic impact on an organization. With GIS data, it can take months and sometimes years to develop an enterprise level operational geodatabase environment – and if this system is lost, it can take several days to restore these data to the last backup. Once the system is rebuilt, lost operations must be repeated to return the database to a current state.

New distributed database technologies promise to expand the way we currently support GIS operational environments. The ability to support a single geodatabase environment through both loosely connected (personal geodatabase checkouts) and distributed database environments provide ways to expand system capacity without breaking up the geodatabase. More loosely connected architecture solutions may significantly reduce the cost and risk associated with current GIS enterprise deployments.

SYSTEM DESIGN PROCESS

System architecture design and platform selection should be based on supporting user workflow needs. The projected user workflow needs, particularly those business workflows supported by dedicated operations staff, must be supported by the selected platform and infrastructure environment. Without a properly balance system architecture design, resources required to support productive GIS operations may not be available.

The platform selection should leverage current IT investments and deployment standards. Platform technology strategies are strategic business decisions, and are directly related to staffing and training investments that support business operations. Operational system availability and security requirements have direct impact on platform selection and hardware costs. The selected solution must also match with available project funding thresholds.

The user requirements analysis should directly support the platform architecture selection, and once that selection is made a system analysis can translate peak user workflow needs to specific peak platform and network traffic loads. Resulting platform and network traffic loads provide a basis for proper system performance capacity sizing and platform selection.

PERFORMANCE SIZING BASICS

Computer technology involves a large quantity of performance variables supported by a constantly changing environment. Enterprise design requirements are seldom understood, and every vendor has his own point of view in addressing performance sizing requirements.

Design consultants must constantly listen to all credible technical views, and piece together a system performance model that can explain and support the valid views of the many experts.

A balanced system architecture design based on validated user workflow requirements establishes a foundation for successful GIS operations. There are many factors that impact performance, not the least being the application design and user workflow. The system design will establish an infrastructure resource threshold that application programmers and system administrators must work within to support productive business operations. Marketing pressures force the performance threshold downward, while technology needs push the same performance threshold upward. The weakest component in the architecture design establishes the system performance threshold. Once the selected system environment is deployed, operations must be supported within the established performance threshold to support user productivity needs.

PLATFORM SIZING MODELS

Platform sizing models help system design engineers simplify and better understand relationships between peak user workflow needs and the resulting system infrastructure process loads. Models provide a method for documenting the views of many technology experts into a system environment compliant with a multitude of performance variables.

Understanding the technology and infrastructure resources required to support use of the technology (platform and network capacity) provides a framework for supporting effective system architecture design consulting services. These models establish a framework for identifying platform specifications required to support identified user workflow needs. Effective use of the models in accurately predicting the performance of real operational environments is validation that the models are correct. Testing the relative performance of new software releases with current model baselines contributes directly to a better understanding of future infrastructure resource requirements.

PLATFORM SIZING TOOLS

Performance sizing tools are maintained to provide a graphical representation of the model relationship between peak workflow platform loads and published platform specifications. Specific vendor platform configurations are represented on the performance sizing chart clearly identifying the peak capacity loads that can be supported by that platform, and the performance capacity relationship provided by available vendor platform alternatives. The performance models are based on published vendor benchmark specifications for each platform environment (SPECint_rate2000 benchmarks support the current vendor baseline). Results are published in the System Design Strategies white paper, published at the following URL: http://www.esri.com/library/whitepapers/pdfs/sysdesig.pdf.

BUILD QUALIFIED STAFF

Once the operational system is deployed, system performance will depend on the staff capabilities to tune, maintain, and make proper use of the system environment. Effective use

of the applications will depend on qualified operations staff. The ability to upgrade and support a current operations environment will depend on the leadership and skills of the GIS support staff. Building a qualified staff is critical to supporting effective GIS operations.

The System Design Strategies white paper provides recommendations for establishing a management team to coordinate and support enterprise operations – executive sponsorship is key to the success of any effective enterprise business operations. The key functional support staff must be identified based on the deployed technology and user workflow needs. Proper training is important to keep staff current with the technology and support administrative requirements. A Project Center Framework tab was recently deployed on http://support.esri.com to connect ESRI customers with information they may need to support a successful and productive project deployment.

MAINTAIN SYSTEM PERFORMANCE

Enterprise GIS operations are supported by complex software environments that must be maintained. Proper testing during initial implementation can significantly reduce deployment risk, and ensure proper infrastructure is available to support user operations. Periodic performance testing can validation proper system tuning, and facilitate maintenance of productive user operations. Database tuning, Web Services tuning, and overall system performance tuning are critical to supporting user performance needs.

SUSTAIN YOUR IT ENVIRONMENT

Technology today is rapidly changing, both in the software available to support GIS operations and the platform and network infrastructure technology available to support GIS operations. Our view of the future is clouded by rapid technology change, and our vision is limited to only a very few years ahead. We no longer have the luxury of deploying an operational solution that does not change.

Managing technology change is the most effective way to support successful business operations. GIS operations evolve over time; the most successful and effective enterprise operations have gown over many years of incremental change – a result of a conscious investment in incremental planning to take optimum advantage of technology change. A current System Architecture Design Strategic Plan for supporting appropriate enterprise evolution is critical to maintaining effective GIS operations. Planning provides a foundation for productive successful business operations.