Preface

Before I was a systems architect, I was a physicist and engineer, and before that, a pilot. I remember more than one mission when I found myself flying in the clouds. What did I do? I learned to trust my instruments. It's very dangerous to fly by the seat of your pants; when the stakes are high, you have to trust your control panel. The stakes are very high when designing a system for technology, too, yet you'd be surprised by how many people think they can feel their way to a proper system design. In contrast to that, I use the performance models in this book to guide my design decisions. As a systems architect, I've created an instrument panel from these models, the Capacity Planning Tool, or CPT (on the DVD accompanying this book), to automate the system design analysis tasks. The CPT releases are currently developed in Windows Office 2007 maintaining compatibility with Office 2003. The CPT will be tested and supported in future Office releases as they become available. These models—developed from years of experience and based on observations of what worked for geographic information system (GIS) users—form the methodology offered to you here, as you seek to mitigate risk and ensure the success of your GIS implementation. The system architecture design methodology in this book has already led to thousands of successful GIS deployments by Esri customers over the last two decades.

This new and expanded second edition further establishes the fundamentals of GIS planning. It is the second book whose mission is to promote more effective GIS planning, especially for large-scale projects, in which planning will significantly improve, if not guarantee, the chances of implementing and maintaining successful GIS enterprise operations. In addition to its focus on the system design process, this book emphasizes integrating planning with system implementation and performance tuning. Geared toward the same broad readership as that of Roger Tomlinson's classic, *Thinking About GIS: Geographic Information System Planning for Managers* (now in its fourth edition), this companion book further bridges the communication gap between those who bring different specialties to the planning effort: CEOs and CIOs who are not technology experts still need to assess what's cost-effective; IT directors and systems administrators need to understand how GIS processes may be more demanding on a system than other software; and GIS managers, specialists, cartographers, and developers may be surprised to learn how many different ways their work factors in to building a high-performance and scalable platform for GIS.

The book is structured to reach all the various stakeholders with the information pertinent to their role. Each chapter begins by emphasizing

the basics in an introduction that, like an executive summary, is intended to make the kind of information decision makers require more accessible to them. Upper managers trust their hands-on people to handle the details, but they still need a grasp on the fundamentals in order to monitor the project and direct those charged with moving it forward to successful completion. Information grows progressively more detailed as you move on through each chapter, with the numbers you need charted in tables and plenty of graphics to illustrate the specific topic or component and its relationship to system design. You can use the CPT on the DVD to support any system design project. Online updates to the CPT are provided on the Esri Press *Building a GIS* companion site at http://esripress.esri.com/bookresources. Also on the DVD are classroom exercises, useful in reviewing the issues of primary importance chapter by chapter, and the *Esri Comprehensive Training Plan for ArcGIS*. The DVD also contains several video presentations demonstrating use of the Capacity Planning tools, and a full eight-hour slide and audio presentation of the *System Architecture Design for GIS* preconference seminar given at the 2010 Esri International User Conference.

As a physicist with a specialty in engineering, I've been putting systems together for 30 years, more than 20 of those as systems consultant and lead technical architect for Esri. This book is a product of both types of experience and describes an approach to system architecture design specifically intended to promote successful GIS operations. With it in hand, you have a time-tested process and a practical tool to help you plan the infrastructure—user requirements, system configuration, platform sizing/selection, and network bandwidth—to sustain the GIS specific to your business needs. My goal is to provide clarity on the fundamentals of system design as they relate to GIS, and specifically to GIS software by Esri, the world leader in developing geographic information systems and my employer. While the latter two are reasons enough for the book's examples to be Esri-software-centric, there's a stronger rationale for this focus. Esri started developing and testing models 20 years ago to figure out how to put systems together to support what our clients do. Working with colleagues, customers, and test teams at Esri, through trial and error and benchmark and performance testing, I began documenting what we were learning in a technical reference document called *System Design Strategies*, and updated it at least twice a year since 1992 in order to share with our customers over the web what we were learning. The *System Design Strategies* technical resource document is now maintained and published on www.esri.com/systemdesign and continues to provide Esri customers with timely online content for GIS planning.

Flying in the clouds was good preparation for a systems engineer starting to work in a software company. When I first got here, users were asking us, "What type of hardware should we get?" and "How should we configure it?" We were fortunate that customers asked these questions early on because it got us thinking about it and testing our ideas in the lab as early as 1991. By 1992, we started using Standard Performance Evaluation Corporation (SPEC) benchmarks, published by the hardware vendors, to adjust for hardware performance changes. (SPEC benchmarks are results of tests set up by the Standard Performance Evaluation Corporation that are used to compare the performance of hardware.) Over the years, this has given us numerical values to reference from the past to compare to present numbers. Since 2000, we've joined with our development staff to conduct controlled performance validation testing to better quantify performance and scalability of each software release. Now, we can demonstrate what we've been talking about for years in terms of what can be done in capacity planning using the models and the CPT in this book. The CPT includes templates for defining enterprise workflow requirements and calculating peak system capacity based on vendor platform selection and system configuration. It can be a tremendous aid in planning and managing a successful GIS implementation—at first launch and at the least cost—when you use it to model your GIS operations. We have added several new tools to the CPT since the publication of the first edition to answer questions on platform capacity, platform sizing methodology, software technology selection, and more.

But a word of caution: Powerful as it is, the CPT is no substitute for understanding the fundamentals of system design, which are also explained in this book. I know from our customers and from students in my classes that it's very tempting to think you can just plug in the numbers and let the CPT do the rest. It is quite a useful tool insofar as you can try out various configurations in theory before buying and implementing the real thing. As a realistic model based on real-world experience, it can get you within the range of the specs you need. But a model is only as good as the numbers you put into it. And when the real world changes, as it inevitably does—we have seen big, rapid changes in the last few decades—you will need to tweak inputs to the model so that it faithfully conforms to the new realities. We need to talk about the real world the models represent, as we do here, because models are just tools. A tool is good, but only in the hands of a person who knows how to use it. And a tool does its best work when used exactly as intended. You wouldn't use a hammer to drive in pilings; you wouldn't start to build without a blueprint, either. You need to know what you need out of the GIS and to describe the information

products that will provide it before you are ready to develop your design solution. This is why a user needs assessment—documenting exactly what your organization requires of a GIS—is prerequisite to completing your system architecture design.

In any cost-benefit analysis, the advantages of modeling a technology solution before committing resources to it are enormous and manifold. And the present time in particular is a strike-while-the-iron-is-hot moment in the evolution of computer technologies. Technologic advancements in software, hardware, operating systems, and the web have recently converged to offer opportunities for more functionality and efficiency at less cost. With the new multicore hardware technology, for example, a dollar buys twice the processing power it bought before. Software offers more functionality and interoperability for the same licensing fees.

We have developed models over time for integration of these advanced and evolving technologies. And now, what we anticipated in the last decade as the next generation of GIS has arrived. The future is here, and not only that, the wherewithal to grow into—and seamlessly adjust to—future changes is here as well, in the scalability of both software and hardware architecture. Future changes will enable more and more GIS sharing of a variety of geographic tools and information layers for use in the new Web 2.0 space. This, in turn, will enable building more and more federated solutions that share pieces of information for use in user workflow collaborations. ArcGIS 10 opens the door for a bright and challenging future with a great many new opportunities. But in order to take best advantage of the opportunities incumbent in all these advancements, you must understand how all these components work together. And you must work together well with colleagues, who bring different but necessary expertise, in order to plan and maintain systems that work well together, too. This book is intended to help you do both.