Getting started with Grid

Clive Harris
Worldwide Strategy Leader, Linux and Grid Services
IBM Global Services

Contents

Executive summary

Despite the growing adoption of Grid technologies, much confusion remains concerning what Grid computing really means, whether any organizations are actually deriving any value today from Grid solutions and whether Grid technology really exists. "Grid explained," a recent white paper from IBM, helps to simplify the complexities of Grid computing by defining Grid terminology, outlining the basics of Grid computing and virtualization and describing how organizations can gain advantage from a Grid solution—both now and in the future. Building on "Grid explained," this white paper helps organizations take the initial steps toward developing a Grid solution and leverages IBM's experience assisting clients with "Getting started with Grid."

Getting started with Grid

Organizations considering Grid computing typically understand the concepts, but want to know the answers to three key questions:

- *Where should we start?*
- *What "roadmap" will provide benefits?*
- *How can we avoid unnecessary risk?*

In response to the first question, the key message is that, with Grid computing, there is not a "one-size-fits-all" approach. Equally true is that, when getting started with Grid, there is not just one roadmap. The benefits can be extensive, but the particular solution for each organization is usually different. However, the solutions typically consist of repeatable patterns and components. This paper highlights common approaches, describes some of the choices involved in selecting an approach, covers key considerations and provides guidance on how to acquire additional help and information.

Organizations are also very focused on deploying proven solutions and therefore typically do not want to take risks with production systems. Some believe that Grid is not yet "ready for prime time." In fact, Grid computing is not new—it has been used in several technical areas for a very long time. Today, many organizations are gaining great benefits from Grid—and increasingly producing measurable business results. The more recent changes include the widespread adoption of standards, the use of Grid for commercial organizations (as well as in scientific and technical areas) and the availability of functional products from IBM and various software partners.

Ultimately, the selection of technologies, tools and products depends on the client's focus and requirements. IBM is particularly experienced in helping clients make the best choices—from various technology, design, architecture, product selection and implementation viewpoints.

Getting results with Grid

IBM has worked with several clients to implement various types of Grid solutions—ranging from the world's largest supercomputers to banks running business applications faster with Grid, from enterprises using Grid technologies to improve efficiencies for horizontal business applications such as SAP to product design shops using Grid to improve engineering collaboration. IBM clients who have already implemented Grid solutions include Charles Schwab, RBC Insurance, Royal Dutch Shell, Nippon Life Insurance, Wachovia and numerous others. Details of these implementations can be found at **ibm.com**/grid. In addition, a growing number of IBM clients are either taking their first steps toward Grid or have completed successful pilots and proofs of concept and wish to apply the benefits of Grid across the enterprise.

The full, widespread business exploitation of all Grid technologies may occur in the future, but many solutions can be adopted for real benefit right now. IBM has solutions and professionals available to help clients achieve benefits early. It is important to know what those solutions are to understand how they may apply to a particular organization and to successfully address any associated risks. This paper summarizes some of the factors to consider as organizations explore which solutions they should adopt.

Introduction

Organizations typically do not consider Grid as a solution in itself. They may be looking for ways to improve the effectiveness of their infrastructure or to improve their business processes through transformation, or perhaps they are seeking opportunities for innovation that will benefit the business. In many of these cases, Grid can certainly play a part (see Figure 1).

IBM has worked with a large number of clients in this area, and has carefully considered the appropriate exploitation of Grid technologies from a customer-centric (business) viewpoint while working with clients to prove the approaches and solutions.

Within this paper, Grid is considered as a wide set of Grid and related technologies such as scheduling, virtualization, provisioning, data, workload balancing and others.

Figure 1: IBM solutions that address customer requirements

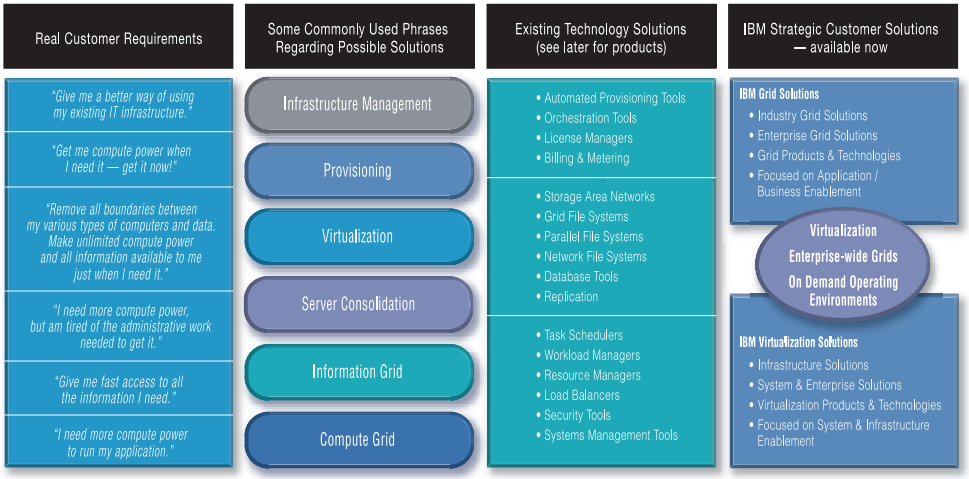

Learn more about Grid

- *See "Perspectives on grid: Grid computing–next-generation distributed computing" at* **ibm.com**/developerworks/grid/library/gr-heritage *for details on how Grid complements and contrasts with other forms of distributed computing.*
- *See the IBM Redpaper "Fundamentals of Grid Computing" at* publib-b.boulder.**ibm.com**/Redbooks.nsf/RedbookAbstracts/REDP3613.html *for discussion material about Grid computing and its concepts, use and architecture.*

Benefits of the Enterprise Grid

IBM refers to the use of Grid technologies or solution building blocks deployed within parts of or across an entire organization as the Enterprise Grid. There are many "point solutions" and specific industry solutions for Grid today, but increasingly clients are looking to exploit these technologies on a wider basis. Typically, the more widely Grid technology is implemented, the more benefits can be gained. Grid is still a relatively new technology within commercial enterprises, so if the full benefits are to be achieved, care and caution are needed as Grid solutions are extended across the enterprise.

The ways in which an organization may initially begin to implement Grid can be influenced by a desire to use Grid across the enterprise at a later date. Some important considerations include selecting the appropriate applications, identifying and involving the right user groups and designing the right technical architecture. Alternatively, an organization may wish to begin by piloting the use of Grid to improve an enterprise-wide business process and its associated applications.

In either case, a stepwise approach—with a long-term view in mind—is usually required. Figure 2 shows a schematic view of how the various aspects of virtualization and Grid can be mapped to a potential timeline. However, as discussed later in this paper, the particular steps for any given client are likely to be unique.

Figure 2: Grid timeline—benefits increase as elements become more complex

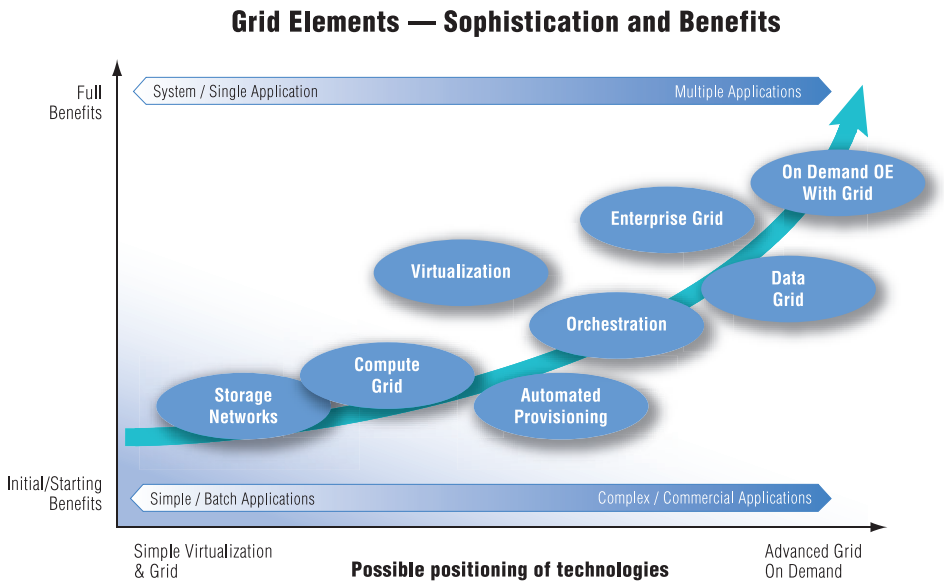

Where to start

The following scenarios describe possible ways of getting started with Grid—by starting with simple, low-risk options and then moving to more sophisticated aspects of Grid if appropriate.

There is not one simple roadmap for Grid, but the following ideas should help in selecting a roadmap based on specific circumstances.

Some key questions to be answered when getting started with Grid are shown in Table 1.

Table 1: Questions to ask when considering a Grid implementation

	Which type of benefits would the organization like to achieve first?	What scope does the organization want for the Grid in the near and medium term?	What technical aspects or elements of Grid should be exploited first?	Does the organization have the necessary skills?
What	• Technical, IT, performance or infrastructure-related benefits • Direct business and application-related benefits	• A department or line of business within the organization • The entire organization (Intra-Grid) • Outside the organization (Extra-Grid or Inter-Grid)	• Compute virtualization • Data virtualization • Application virtualization • Service virtualization	• Technical skills • Business and process skills
How	• Select an initial application that has direct business benefits • If that is not possible or appropriate, select a business or application area that would have the widest positive impact	• Select an initial scope that has an enterprise-wide aspect to it as early as possible	• Determine the technical elements to be implemented based on the business objectives • Define an architecture that allows the organization to incorporate each of the future Grid elements at the appropriate time	• Take a stepwise, business-needs approach, which is always preferable to a "big bang" approach
Why	• This often makes the justification for the Grid solution easier and may help to get wider "buy-in" to Grid across the organization	• This can help address organizational concerns as well as technical concerns, and can help the wider organization see the benefits of Grid • Early return on investment quantification and demonstration (through piloting) may be beneficial	• Do not implement elements of the Enterprise Grid that are not yet needed	• This allows in-house technical skills to be acquired at the same time

Possible steps to virtualization and Grid

This section describes some common ways in which IBM has helped its clients deploy and extend their Grid solutions. This section focuses on the logical and conceptual steps taken; appropriate technologies, products and technical solutions that allow clients to achieve these steps are described later.

Organizations are often likely to start with *system*-related virtualization as they make their first moves toward Grid technologies. System virtualization is also often likely to be an enabler for the more business- and application-focused benefits that Grid solutions can provide. It should be stressed that the concepts here are not new, but are now being adopted at a greater rate.

Initial steps, as shown in Figure 3, are first to virtualize similar resources from a system perspective, and then to do the same with unlike resources, or *heterogeneous systems*. This allows more rapid deployment of application and Grid solutions to exploit the system virtualization.

Figure 3: Initial steps for virtualization and Grid

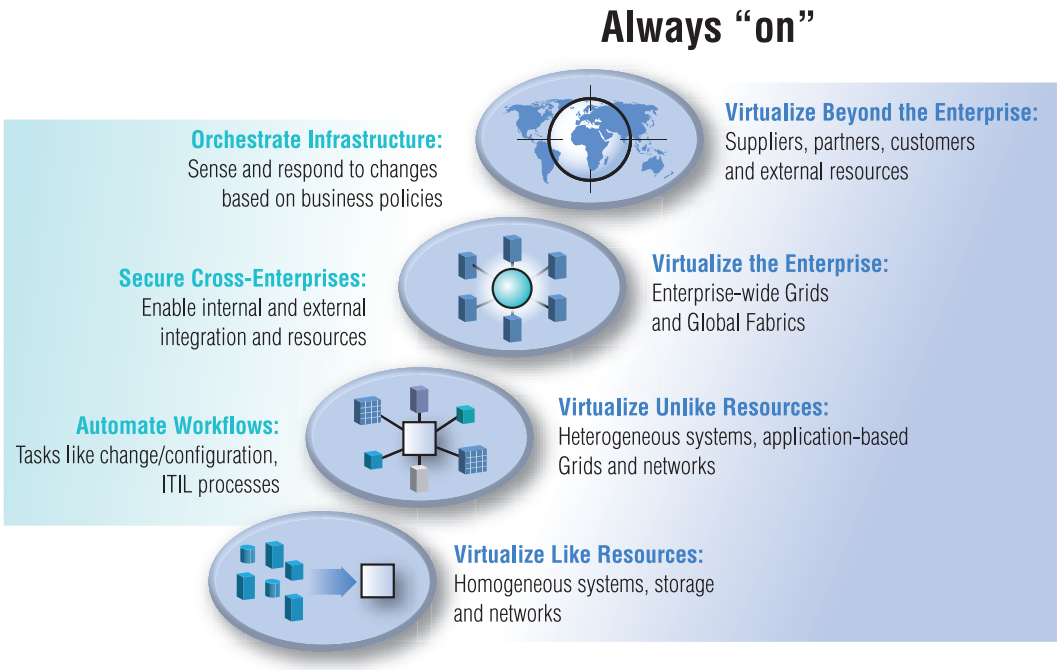

Virtualization of system resources can potentially provide substantial benefits, and is achievable today with a range of IBM and IBM Business Partner products. One such product example is IDI, which is a chargeable component of the IBM Virtualization Engine™ technology. This particularly applies to the many enterprises that run SAP software.

IDI allows the organization to build a pool of shared server resources and to share these resources between SAP application servers across various SAP systems.

It also supports the specification of application service level agreements (SLAs)– for example, 80 percent of dialog steps must be completed in less than 1 second. IDI monitors SLA conformance, automatically assigns resources from the shared pool if SLAs are violated and automatically moves resources back to the free pool if SLAs are overachieved. IDI also provides accounting information about resource usage.

Shared resource pools can contain resources from different platforms. For example, Linux® and AIX® platforms and Linux on IBM @server® zSeries® servers can be provisioned to an SAP system depending on availability and customer's business goals (see Figure 4).

This type of virtualization, in this case for SAP, can potentially provide significant cost savings for many organizations because of the increased utilization and reduced number of systems needed. It also allows for faster deployment of new SAP systems.

Figure 4: An example of virtualization with SAP

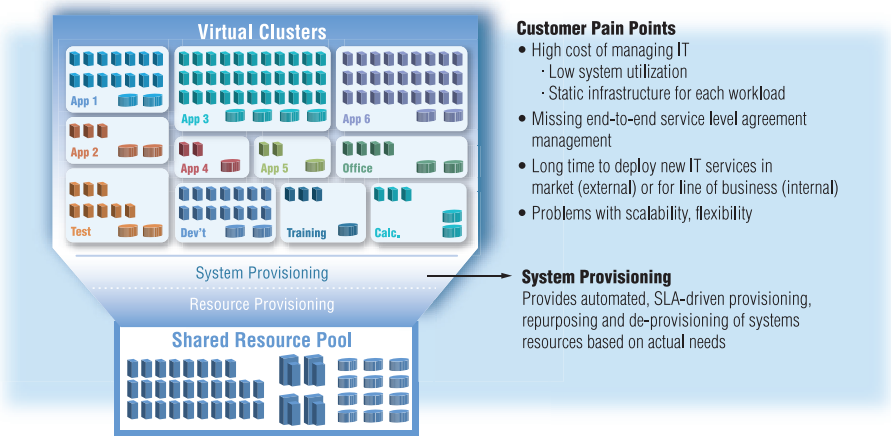

IT resources are pooled, virtualized and allocated dynamically to satisfy changing business and workload needs.

Some possible steps to Grid

By looking at the Grid and application exploitation steps, organizations can consider a similar possible stepwise approach that minimizes risk and allows for early benefits to be achieved.

Figure 5: A stepwise approach to implementing Grid

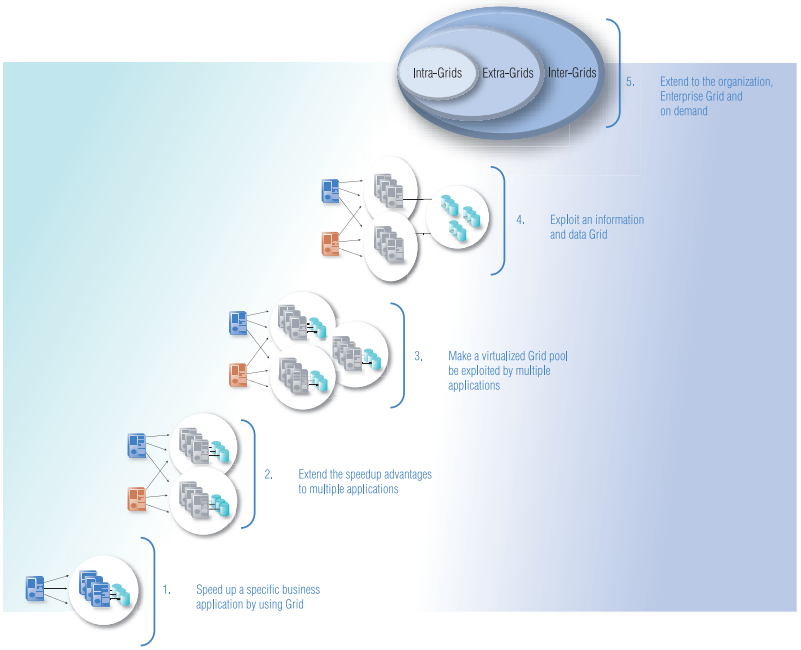

As shown in Figure 5, these steps may include:

1. Step 1: Speed up a business application by using Grid
2. Step 2: Extend the speedup advantages to multiple applications
3. Step 3: Create a virtualized Grid pool that can be exploited by multiple applications
4. Step 4: Exploit an information and data Grid
5. Step 5: Extend to the organization, Enterprise Grid and on demand
6. Other possible steps in the future

Each of these potential steps is described in the following sections.

Step 1: Speed up a business application by using Grid

The first step is to take an existing application and break it up so it can leverage multiple compute resources simultaneously (see Figure 6). This can improve the performance of the application and allow an organization to gain better use of existing servers and infrastructure.

There are many ways of achieving this even for applications that are not designed to run in parallel. Such examples could be particular business applications for a specific department, or they could be enterprise-wide applications. Either would be transparent to the user, who would get the same application and user interface but with an order of magnitude run-time improvement.

The application in question may take advantage of spare capacity on other existing servers, or may simply use servers that are not always in use and can be used for this purpose.

Figure 6: Breaking up an application so that it can use multiple compute resources simultaneously

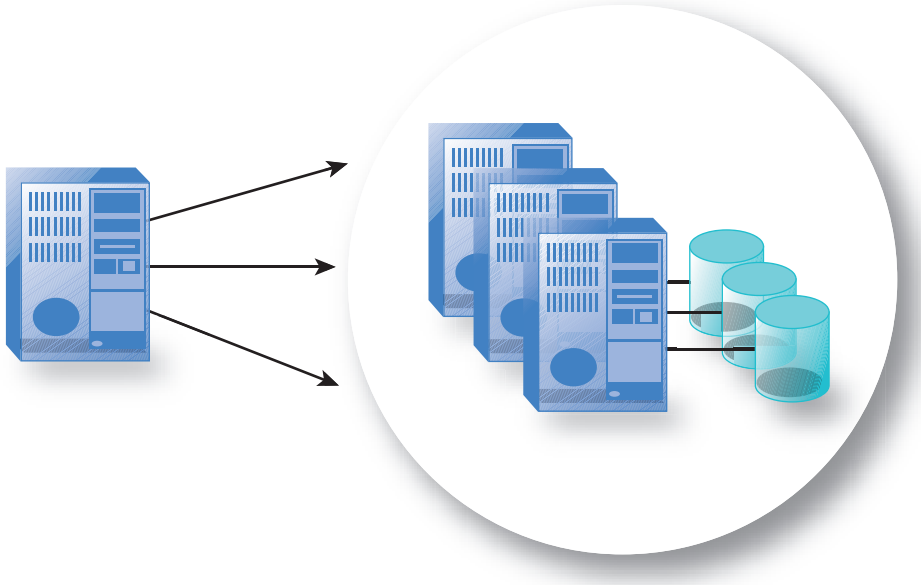

Examples of business applications that can take advantage of this approach are clash analysis (in the automotive industry) or analytics acceleration (in the financial services industry).

Examples of enterprise-wide exploitation include compile Grids for the technical and development community across an enterprise and anti-virus Grids to better control and respond to anti-virus checking across the enterprise. Yet another example would be to use Grid in an enterprise-wide manner and apply it to the many aspects of the typical SAP environments that are in use across the organization.

Step 2: Extend the speedup advantages to multiple applications

Having achieved success with one application and tested the Grid technologies, organizations often extend this success to other areas. Adding another application to the scope of the Grid is an obvious way of moving forward (see Figure 7).

This second application may be owned by the same department or line of business, it may be owned by another department or it may be an enterprise application.

If the intent is to select an application that crosses organizational or departmental boundaries, that process should begin early. The obstacles to Grid are not just technical in nature, but also can be organizational or cultural. It is helpful to overcome these obstacles sooner rather than later if the full benefits of Grid are to be achieved early.

Adding more applications to the Grid environment can provide greater benefits. Careful selection of the right applications allows the applications to "complement" each other—perhaps having workload peaks occur at different times or requiring different types of resources that can be shared very effectively.

Figure 7: Extending Grid capabilities to additional applications

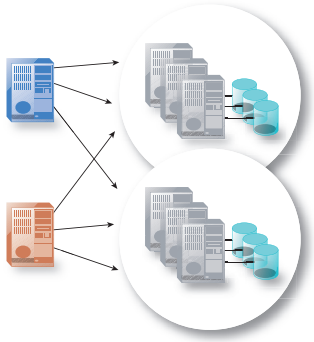

Step 3: Create a virtualized Grid pool that can be exploited by multiple applications

As more departments or applications are added, the Grid environment will be well on the way to a "virtual pool" of servers that can be used and exploited across the enterprise. At this stage, it may make sense to specifically start creating such virtual pools of servers (see Figure 8), which could include underutilized servers or servers utilized only at certain times. These servers could be made available for other applications and workloads.

There will always be some servers and applications that require dedicated processing, which may be important to some organizations from a security viewpoint. However, the more an enterprise can take advantage of shared pools of servers and more dynamic application allocation, provisioning, automation and workload balancing, the more benefits that enterprise can reap.

Figure 8: Building a virtual pool of resources

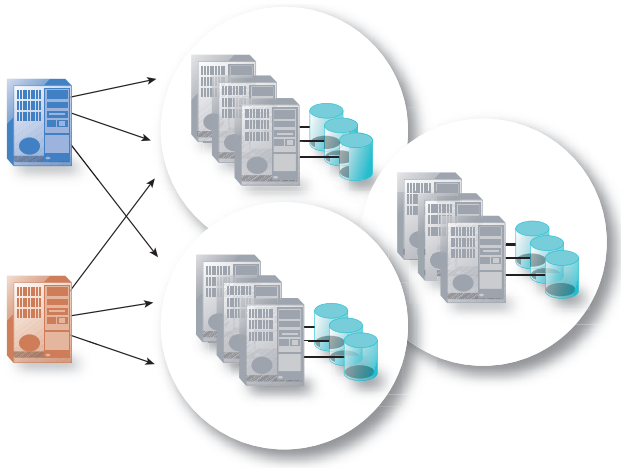

Step 4: Exploit an information and data Grid

In this step, Grid concepts are applied to information and data rather than just to applications or compute power (see Figure 9). This step may be reached earlier than outlined in this paper; in fact, for some clients it is the first step.

Data virtualization can be achieved in various ways and for a variety of purposes. It may be that a data or storage Grid is used to access data that is not currently available across the entire enterprise. Alternatively, a data Grid may be implemented to obtain a significantly higher level of performance for data and information access.

Figure 9: Applying Grid to information and data resources

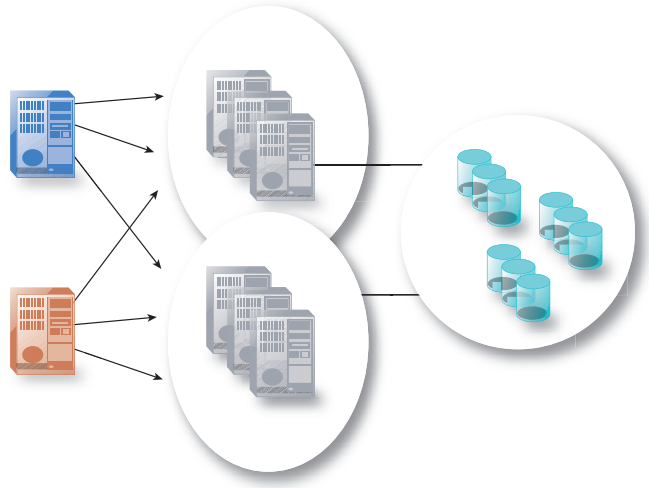

Step 5: Extend to the organization, Enterprise Grid and on demand

Having achieved the previously discussed steps or similar approaches allows an organization to continue extending Grid technologies to gain more and more enterprise-wide benefits (see Figure 10).

An Enterprise Grid is simply the exploitation of standard Grid technologies and solutions, but within and across the entire enterprise. Once an Enterprise Grid is achieved, organizations have an operating environment that supports an on demand business.

The IBM on demand operating environment defines a set of integration and infrastructure management capabilities that an organization can use—in an incremental and modular fashion—to become an on demand business. Three characteristics differentiate the IBM on demand operating environment: infrastructure management, automation and virtualization.

Figure 10: Extending Grid across the enterprise

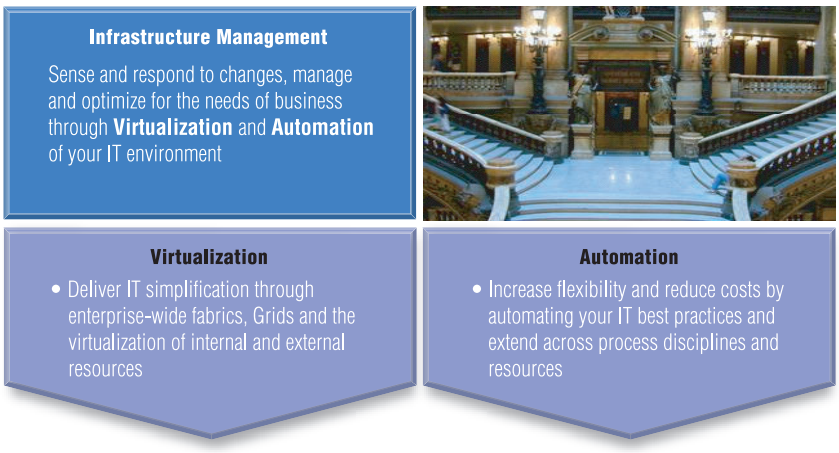

Other possible steps in the future

Other steps that may be considered in the future would include extending Grid beyond the existing sites and even perhaps beyond the existing organization (see Figure 11).

Many other factors come into play if such strategies are planned. Security and network bandwidth are obvious factors. However, there are advantages to be gained by having Grid solutions spanning multiple sites or even continents if the organization is multinational. Around-the-clock workload balancing could provide significant benefits, and the options for disaster recovery and high availability would be greatly enhanced.

Not many commercial organizations are ready to deploy Grid solutions that span into external organizations, but this is common in the scientific and technical community where some of the application characteristics are different. However, in the future, more widespread adoption of Extra-Grid (extranet) and Inter-Grid (Internet) solutions is likely.

Figure 11: Extending Grid beyond the enterprise

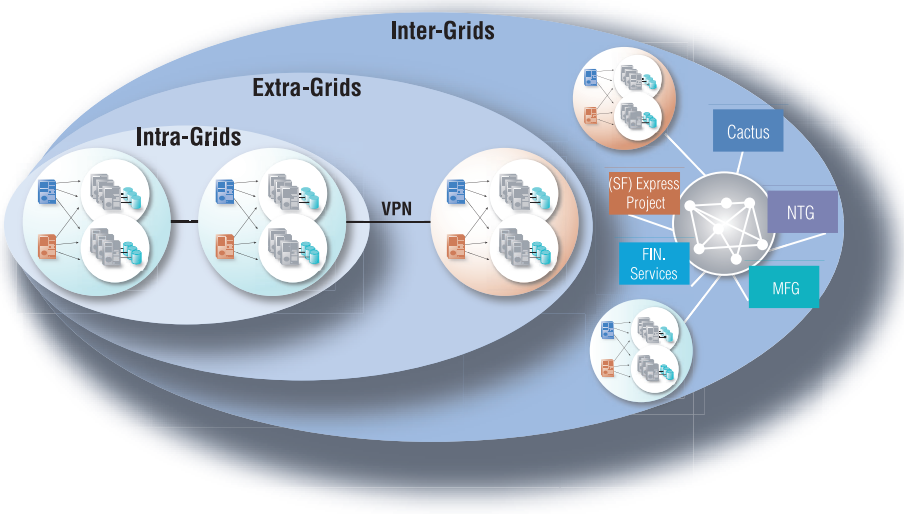

Technology and product choices

As mentioned earlier, Grid is not an off-the-shelf solution. Organizations cannot just "buy" a Grid and get the software delivered in one box. However, there are several products and solutions from IBM and IBM Business Partners that can provide an organization with a set of building blocks to enable rapid initial progress in building a Grid. In the past, some of the more technically minded clients created and developed their own Grid solutions.

IBM recommends that clients purchase standard products and tools to build Grid solutions. Such solutions are more likely to be supported by third parties as well as more stable and maintainable. One of the key reasons for this being the favored approach to implementing Grid today is the continued maturation and adoption of Grid standards. Standards are an important consideration in Grid solutions because Grids often contain many components. Therefore, standards not only facilitate interoperability, but they also allow organizations to assemble Grids from standard components.

The key technology and product choices usually involve Grid middleware and tools that provide the Grid functionality. Some fundamental technologies are typically part of any complete Grid solution. IBM, working with its Business Partners, is uniquely positioned to provide all of the pieces of a complete Grid solution, rather than just limited scope solutions.

The key technologies and requirements that comprise a comprehensive Grid solution consist of capabilities in the following functional areas:

- *Scheduling*
- *Resource management*
- *Provisioning*
- *Data virtualization*
- *Full integration of the above areas*

As clients increase their usage of Grid, capabilities from across these areas are functionally integrated to successively increase the sophistication of the overall Grid architecture and solution. For example, clients may begin with a compute Grid based on scheduling, but as they increase the number of applications and users, integrating provisioning with scheduling adds significant value to the overall solution.

Scheduling

IBM provides its own software products and also works closely with the key Grid independent software vendors (ISVs) today to provide comprehensive scheduling functionality. In the future, customers may wish to have enterprise workload scheduling that encompasses very flexible enterprise-wide workload and resource brokering. In an enterprise setting, clients may want to support a distributed network of schedulers working together to coordinate workloads across multiple environments such as clusters. In a large organization, it is unlikely that one scheduling solution can meet the needs of all business units and users, and therefore scheduling standards become an important consideration.

Traditionally scheduling has been segmented into three categories: static scheduling, dynamic scheduling and cluster scheduling. As Grid adoption matures, being able to unify various scheduling environments (static, dynamic and cluster) across an enterprise into a common and consistent view allows an organization to view scheduling as one discipline.

Today IBM offers Tivoli® Workload Scheduler, WebSphere® XD and LoadLeveler® software as scheduling solutions. IBM also partners closely with DataSynapse, United Devices and Platform Computing to integrate their scheduling products into the IBM Grid solutions.

Resource management

Resource management is vitally important to Grid because everything in Grid computing and virtualization revolves around resources. Work is scheduled to resources. Entities are provisioned to resources. This is not necessarily a new concept or one that is specific to Grid computing, but excellence in this area is a necessity because of the ubiquity of resources in the Grid world.

If scheduling moves work to resources, dynamic resource management or orchestration moves resources to the work. Scheduling is the demand side of the equation, and orchestration is the supply side; they are symbiotic.

IBM products in this space include Tivoli Intelligent Orchestrator and e-WLM within IBM Virtualization Engine.

Provisioning

Closely tied to resource management is provisioning. Many types of logical resources need to be provisioned across Grid infrastructures in a highly dynamic fashion. Software environments need to be provisioned, licenses need to be provisioned and data needs to be provisioned. Therefore, provisioning is a defining technical characteristic of Grid environments. IBM's core product in this area is Tivoli Provisioning Manager.

Data virtualization

Grid applications require data access across the distributed system in a manner that provides a variety of interfaces, federation of different data sources, location transparency, global data naming and high-performance data access. This is referred to as data virtualization. True data virtualization is challenging because of the wide variety of data formats, access protocols and quality of service requirements. But the potential benefits to organizations are significant. By unifying information access in a potentially seamless and transparent manner, organizations can address their motivations to enhance employee productivity, respond to variable demands and quickly adapt to changing business requirements.

IBM products that support this type of functionality and that are often used in Grid environments include IBM DB2® Information Integrator (DB2 II), Global Parallel File System (GPFS), Network File System (NFS), SAN (storage area network) File System (SAN FS) and SAN Volume Controller (SAN VC).

Integration of these technologies

What is perhaps clear at this stage is that many Grid solutions are built using standard IBM and IBM Business Partner products that are also used in non-Grid environments. For this reason, IBM has been able to deliver Grid solutions to its clients rapidly—because the Grid solutions are usually built using tested products as building blocks.

However, IBM is unique in its ability to integrate the various technologies into enterprise-ready and robust Grid solutions. IBM's experience in this arena is second to none, and IBM has Grid solutions and combinations of products and technologies that have been deployed with many customers. IBM also

supports a dedicated Grid Integration Center in Austin, Texas, to validate and configure the integration of IBM and IBM Business Partner products in Grid solutions. By prefabricating Grid infrastructure solutions, clients may be assured of the overall interoperability of a Grid solution.

IBM Grid Reference Architecture

In addition, IBM has defined a Reference Architecture that it uses in the design and deployment of Grid solutions for clients. A summary of the major components of this Reference Architecture is shown in Figure 12.

Figure 12: IBM Grid Reference Architecture

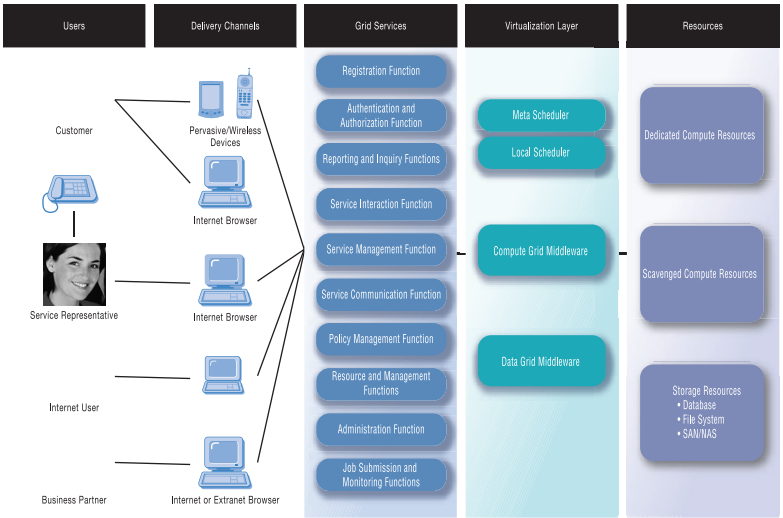

Products from IBM and others can provide each element of the functionality; of course, not all of the functionality is required initially. IBM helps clients to design and develop suitable, appropriate Grid solutions using this approach, and can also match the requirements to the most suitable products. Typically, it does not make sense to simply select a set of Grid middleware products without careful consideration of the requirements first.

A typical technical Grid solution

Putting each of the pieces together and considering the "end game" that may be achieved after all of the steps previously described, a typical solution may look like the one shown in Figure 13. This solution is one that has actually been built and is being used today. The solution uses WebSphere Grid Portal, a partner Grid Scheduling product, Tivoli Provisioning Manager, Tivoli License Manager, SAN FS and other technologies and products.

What becomes clear as such solutions are constructed is that a Grid solution consists of several essential elements, some of which will become more important as Grid is used across the enterprise. It is possible to build simple Grid solutions that comprise only some of these elements, but it is important to design solutions that allow the Grid to grow and be more widely exploited. There are known cases where deployed Grid solutions have not been flexible, expandable and scalable enough to satisfy future growth and requirements.

IBM recommends that clients define the most appropriate Grid architecture, and then implement and deploy just those pieces that are needed initially and can provide the most benefit. But a defined "blueprint" helps ensure that the solution can grow within the requirements over time. IBM is well positioned to help customers to define and implement such Grid solutions.

Figure 13: An example Grid deployment

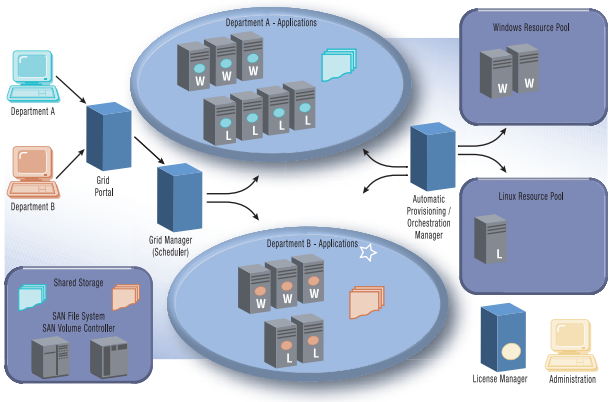

Creating a Grid roadmap

IBM can help clients to create a client-specific roadmap for Grid. IBM has skilled and experienced professionals, specific tools and standard services that can help clients to define such a roadmap. IBM also uses a repeatable process for building such solutions—the IBM Global Services Method. This method allows IBM to reuse solution components around the world and to help ensure optimum roadmaps and designs for Grid solutions.

Specifically, IBM can assist in the selection of applications, consult on which business processes need to be transformed, build the justification and take part or full responsibility in the design and implementation of the Grid environment.

Defining a roadmap to Grid, such as the one shown in Figure 14, can be completed very rapidly. IBM's experience with other clients can help ensure that pitfalls are avoided and appropriate technical solutions are defined.

Figure 14: A roadmap to Grid

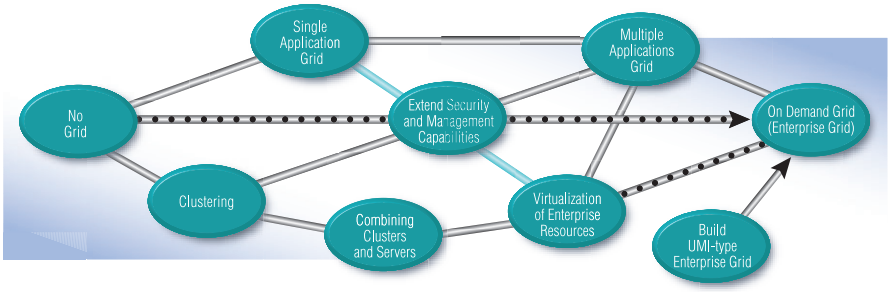

Defining a roadmap with assistance from IBM

IBM regularly works with clients to help them define the most appropriate ways in which Grid can help overcome pain points—now and in the future. A common approach during joint workshop sessions is to consider immediate business or technical pain points, as well as strategic pain points and to construct a current and desired capabilities matrix, as shown in Figure 15.

Figure 15: Current and desired capabilities matrix

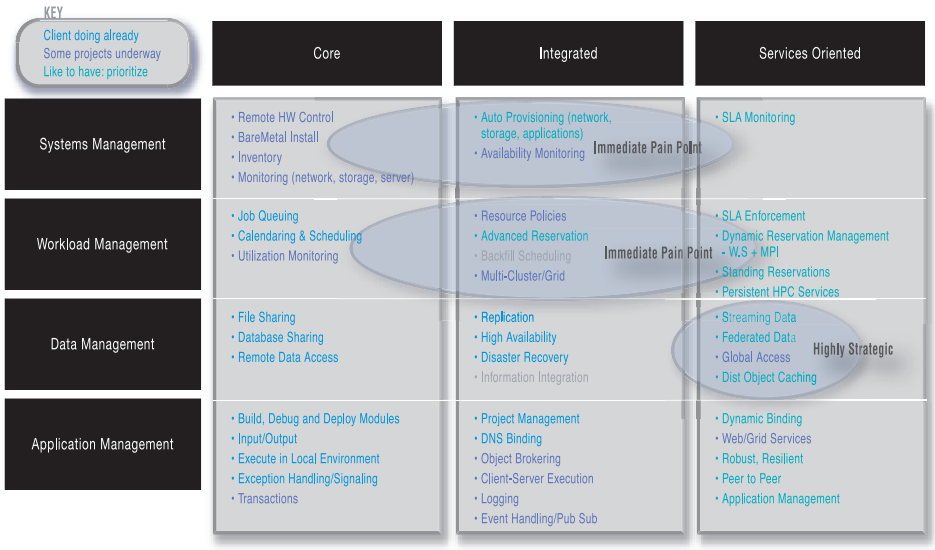

Inhibitors to Grid and overcoming obstacles

As previously mentioned, Grid is still a relatively new technology, so there are almost certainly obstacles to be overcome as Grid solutions are developed and deployed. Having worked with many clients, IBM understands these obstacles very well, and in most cases, has developed solutions and remedies to these inhibitors.

Most of the inhibitors can be considered as either:

1. Technical inhibitors—requiring a technical solution
2. Political, cultural and non-technical inhibitors—still requiring a solution, but not technical in nature
3. Perception and awareness—often overcome once the aspects of Grid are fully understood

Figure 16 shows the major obstacles to Grid adoption, as determined from a recent study of IBM clients. In most cases, IBM has developed solutions to these problems. A series of white papers from IBM addresses each of these aspects and provides further information and guidance to clients. For example:

- *IBM has a range of advanced tools to assist with defining business value, return on investment and total cost of ownership and with creating a business case for Grid. These tools and approaches are described in the IBM white paper entitled "The business value of Grid"*
- *The IBM white paper entitled "Grid explained" discusses terminology relating to Grid and the associated technologies*
- *IBM has published a series of customer briefs that describe how clients have already used Grid to their advantage*
- *The IBM white paper entitled "Six strategies for grid application enablement" describes the various ways in which clients can take advantage of Grid without having parallel applications*

Figure 16: Major inhibitors to successful Grid deployments

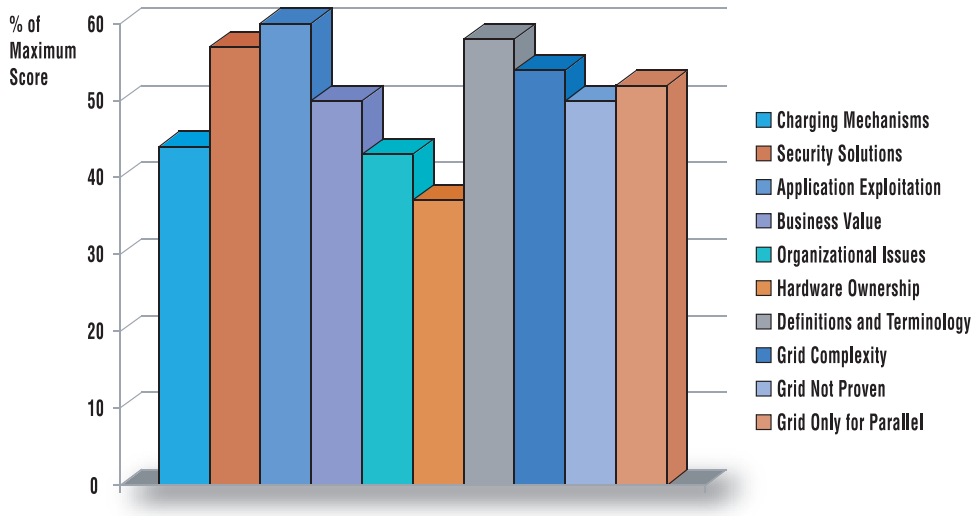

Assistance and support from IBM

By leveraging its extensive Grid experience (internally in IBM and with clients), IBM has developed a range of support and services for clients who wish to adopt Grid solutions. IBM has specific assets that can provide clients with a timely and efficient outcome.

In particular, IBM can provide client case studies and references, can help with several early aspects of a Grid project such as creating a business case and can offer a full set of services to deliver a Grid solution such as solution design and implementation.

A variety of pre-sales support and advice is available from the dedicated IBM Grid teams around the world. The full range of IBM Global Services offerings that are applicable to Grid are shown in Figure 17, but certain services may be more applicable than others. Such an example is IBM's Grid Value @ Work, which is shown in Figure 18.

Figure 17: IBM's Grid Value @ Work

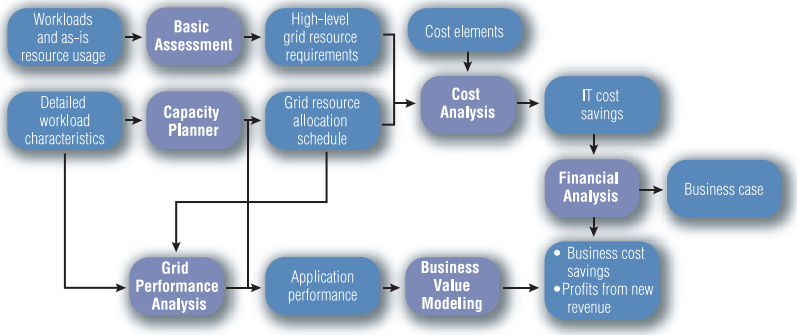

Figure 18: Catalogue of Grid services from IBM

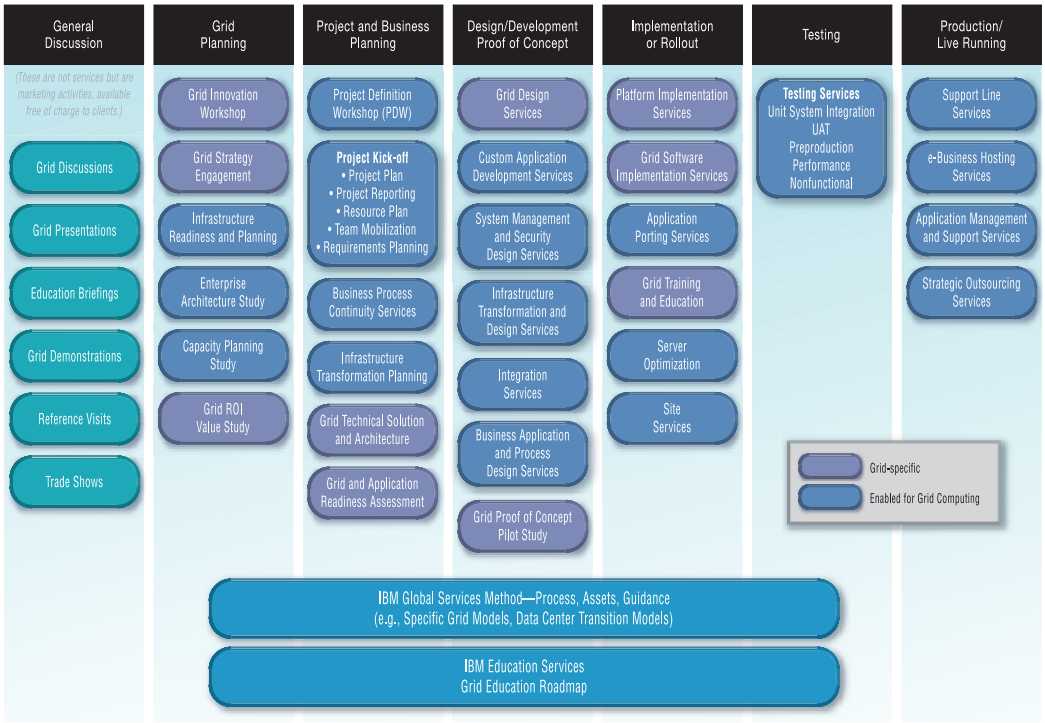

Appendix: Pointers to further information from IBM

- *Information Grid Part 1–information Grid infrastructure*
  - **ibm.com**/developerworks/grid/library/gr-info1
- *Information Grid Part 2–Blueprints to build Grid virtualization layers*
  - **ibm.com**/developerworks/grid/library/gr-info2
- *Information Grid Part 3–information Grid scenarios*
  - **ibm.com**/developerworks/grid/library/gr-info3
- *Grid frequently asked questions*
  - **ibm.com**/grid/about_grid/faq.shtml
- *An overview of Grid: "Grid computing: What are the key components?"*
  - **ibm.com**/developerworks/grid/library/gr-overview
  - *Offers a high-level overview of Grid computing concepts and the components involved, and discusses the Globus Project*
- *Enabling applications for Grid: "Six strategies for grid application enablement"*
  - **ibm.com**/developerworks/grid/library/gr-enable
  - *Provides a model for applications that can exploit the Grid at different levels*
- *Application design: "How grid infrastructure affects application design"*
  - **ibm.com**/developerworks/grid/library/gr-infra.html
  - *Provides a primer on considering both the Grid structure and the type of application when crafting or adapting an application for a Grid*

IBM Corporation
Systems and Technology Group
Route 100
Somers, NY 10589
U.S.A.

Printed in the United States of America
January 2005

GRW00899-USEN-00