THE CLIPPER GROUP
Navigator™

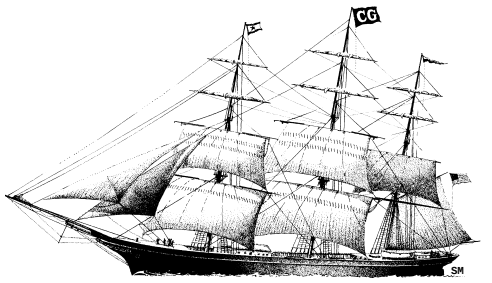

Navigating Information Technology Horizons℠

Published Since 1993 Report #TCG2006086 September 30, 2006

IBM Multiplexes the Data Center with SOA and Grid

Analyst: Anne MacFarland

Management Summary

Business IT environments continue to struggle with change. Processes are used by many kinds of users, and need to be exposed to their various populations in carefully masked sets of functionality. Some applications are in constant use, while others are used only at the beginning of projects, or at the end of the fiscal year. Most applications need more and different resources at some times than at others. The urgency that drives data center use is a matter of business demand, and cannot always be scheduled.

These characteristics are not well matched by the asset orientation that pervades the way IT organizations traditionally have deployed, integrated, and managed IT assets to deliver applications. The low-asset utilization and licensing charge overages are painfully-obvious evidence of this mismatch, but there is more. The bigger and more comprehensive applications are, the harder it is to integrate or amend them safely. This is not a matter of bad practices, but of the heritage, based on the expense and limitations of computers, of focusing on the most expensive assets – servers and storage.

These days, the most expensive asset is time. By focusing on the time dimension of functionality (the service) rather than the physical dimension of functionality (servers and storage), one exposes a set of metrics that allows the data center to better align operations more closely with business requirements. Think of the changes that the information provided by on-line tracking has made to the shipping industry. By adding bar codes to packages and devices to document their location, the shipping industry entered a new stage of profitability. Additionally, multiple qualities of service could be supported. Instead of *as soon as we can*, we got *next day*, *three days*, and *economy* options with rates to match. Customer expectations and satisfaction grew. But, one often-ignored aspect is how much tracking improved the processes for the shipper. Fleets of vehicles could be used more effectively. Routes could be optimized before the packages arrive at their final distribution point.

XML is like the shipper's bar code. Its use lets IT environments work in a more sensible way now that time is of the essence. It allows software to be written, integrated, and evolved into the very big but very agile software environments that on-demand business requires. **When used to manage physical assets, this descriptive approach is embodied as Grid**[1].

IN THIS ISSUE

[1] Grid is built around the idea of implementing OS-style scheduling and control systems over distributed assets, using discovery, self-advertisement, and a services registry to optimize for whatever the workloads require (often time to completion, number of iterations, or simply being able to process very large calculations). Grid architectures use XML to articulate processes and policies, allowing automation of standard IT routines and escalation of exceptions. This allows IT staff to get away from administrative routines and instead focus on business issues.

When used to allow software to be written and evolved as a series of intelligent, self-describing, aggregated modules, it is known as Service Oriented Architecture, or SOA[2]. In both cases, the move from hard-coded integration to a more loose affiliation governed by requirements (policies and governance) and discoverability (fostered by registries and the descriptive nature of XML) will break the draconian correlation between the size of the system and the time it takes to evolve it. This will allow the Data Center to embrace change as initiation of a new service on existing infrastructure – a matter of negotiation, not a powered-down reconfiguration.

With change a normal process, the enterprise becomes less constrained in its evolution, and can be more innovative in how it addresses markets. From an operational perspective, applications can be more closely tailored (and re-tailored) to the business. The self-descriptive nature of the process flow allows for the use of accelerants (be they scale-out scale-up, or offload architectures) to be well targeted.

Grid's efficiencies, like the time-division multiplexing over a telephone wire, make more efficient use of assets and enable swap in/swap out of assets via a service redeployment as a natural part of operations. SOA's modular application composition allows effective business processes, targeted to particular addressable opportunities, to be built. In beer-advertisement terms, SOA is the "tastes great," and Grid is the "less filling." Of course, you want both.

IBM has products, services, and a lot of experience in both SOA and Grid. It has a strong support for the open standards that underlie both architectures, having helped develop many of them. It has middleware-based control software, built on common elements, to support both SOA and Grid environments with a minimum of semantics arbitration. For more details on IBM's complete SOA/Grid story, please read on.

The Focus on Service

Much has been said about the particulars of Grid and SOA. However, **the importance of Grid and SOA lies not in their particulars but in the way that service-based architectures reposition the data center from a focus on often magnificent, but rapidly-aging technology monuments, to a focus on the services created and the infrastructure that can best support them.** This service orientation, instantiated in both Grid and SOA, has some key characteristics.

The Central Role of Virtualization

Service, as a concept, is a virtualization of all the things needed to accomplish it. A service is offered by describing what it will accomplish, not all the things it will use to accomplish it. This behind-the-scenes handling of administrative details underlies the familiar concepts of processing and storage virtualization. Grid has been used to embody processor and/or data virtualization, but **what is really needed for full data center efficiency is virtualization of all infrastructure components – including network elements and workloads – something better named *infrastructure virtualization*.** Then, when SOA brings together application elements in a runtime, they can be deployed wherever can give a quality of service at the least cost.

A Common Set of Consistency Points

Both infrastructure virtualization and SOA require consistency in operations, lately embodied by ITIL[3], but also growing out of development and testing organizations of various sorts.[4] Industry standards, such as XML, create the opportunity for more kinds of consistency initiatives, like that for *Common Event Management*. Such initiatives build more points of commonality that let administrators construct policy, which is then implemented by the system. The consistency of

[2] SOA, Service Oriented Architecture, is a software cousin of Grid, not a clone. Its focus is on deconstructing large brittle applications into more-easily-evolved components, which can be linked in runtimes to support, in particular, customer-facing applications. This component approach allows applications and information to be aggregated and integrated by policy rather than by the hand crafting of traditional enterprise application integration (EAI).

[3] ITIL stands for Information Technology Infrastructure Library, and is a set of best practices for data center operations.

[4] One example is the open-source oriented SOALink community, briefly described in **The Clipper Group Navigator** dated May 7, 2006, entitled *SOALink – A Seedbed for SOA Growth*, and available at http://www.clipper.com/research/TCG2006033.pdf.

coordination possible through consistent use of XML gives an opportunity to build another layer of consistency – *semantic consistency*, which becomes a core component of the terms and conditions of service contracts.

Support for Heterogeneity

The service orientation of both virtualized infrastructures and SOA supports the long-term evolution of the data center, as contrasted with the swap-out of assets and erosion of value every three-to-five years.[5] This prudent longer-term continuity inherently is heterogeneous in nature, because products change. The support for heterogeneity lets data centers leverage whatever hot new technology fits their business model. It also promotes independence from any one vendor or platform and gets enterprises above the hype and venom of protocol wars.

Support for heterogeneity is not a matter of "open" versus "proprietary", an issue that often runs in close parallel with industry standards, but should not be confused with it. **Open versus proprietary is a matter of code and its availability to independent developers.** The more a developer or systems integrator knows about an application's code, the more effective addition and integration of new features can be – if you're doing it at the code level. **However, for all the elegance you get in a code-to-code integration, you lose *reusability* and, in the long run, *time*.** In business, timeliness is often a better predictor of success than internal elegance. Arbitrating services is more efficient than endlessly reworking code.[6]

Common Capabilities Needed to Civilize Distributed Systems

The worlds of hardware and software are very different, and SOA and the infrastructure virtualization of Grid are no exception. However, they do use many of the same tools and techniques. They use many of the same XML variants, such as *WSDL* and *BPEL*. Both use the ideas of registries, repositories, and brokers to discover entities and negotiate services between them, but this is as basic in rapidly evolving environments as humans' drinking water. In SOA, the brokering is in composing the applications that are to run. In Grid, it is used to schedule and orchestrate the deployment of the applications on the hardware that will supply processing and connectivity services. Load balancing helps in both environments. In both, some precipitated form of identity must be used for identification. Self-describing data entities are also needed to keep the system rational and secure. Reconciliation of inconsistencies – be they of policies or of semantics – must be built into the system of both SOA and Grid.

Because they share a common approach to common problems, there are synergies to be gained by using both architectures. A greater articulation of application flexibility to address the needs of particular users is gained by targeting the application modules to those users, and by running the application on an infrastructure to meet user needs for immediacy of response. The ability to scale out at a moments notice, and to treat failure as an incapacity, to be addressed by substitute capacity, is huge. **Nevertheless, the biggest synergy, achieved by using SOA and full infrastructure virtualization, is *greater transparency of the technology infrastructures* that underlie the business – an ability supported by the RAS[7] features of the hardware and the runtime controls of the application.**

What IBM Brings to Bear on these Challenges

Probably centuries of person-hours of experience[8] underlie the Grid and SOA products and services that IBM brings to market. That experience, and IBM's increasing abilities to expose and leverage it, benefit customers that are planning broad implementations because, as anyone who has had to deal with badly managed services will attest, comprehensive scope and good quality are the primary *make or break* attributes of service systems. More

[5] This old approach becomes increasingly spendthrift, as assets get more reliable.

[6] This is not to say that code does not matter – code, and deftness in manipulating the code, is key to building new technological capabilities of capture, analysis, and presentation –capabilities that underlie the human concept of progress.

[7] Reliability, Availability, Scalability.

[8] In part, this is because the time-based multiplexing that both SOA and Grid leverage is a key heritage of mainframe architecture. The challenges of such structures are well documented at IBM

importantly, if you cannot provide the quality consistently, pretty soon you won't be in the business of providing it at all.

IBM has the necessary breadth of domain to have experience and products to address the challenges posed by SOA and Grid. Its approach to virtualization, to identity, and to management are all platform-agnostic. A common approach to virtualization provides more consistency across rapidly changing software and hardware environments. Moreover, pervasive aggregation of usage measurement[9], charge-back mechanisms, and billing is another requisite component of these service environments, particularly if they extend between organizations. If charge-back is coarse-grained or based on a percentage of the whole, someone is probably paying too much – and the data center will be subject to the harassments of commoditization mentioned earlier.

IBM's consistent use of standard interfaces, common points of consistency, and collaboration on many fronts[10] give it a great deal of credibility. Its *Grid and Grow* and *WebSphere* SOA elements benefit from both the breadth of IBM platform coverage and the well-regulated common elements that are the basis of IBM Software – including *Tivoli*, *DB2*, and *Lotus*.

Conclusion

Service-oriented datacenter practices - in the form of Grid and SOA architecture - are a long-term boon to businesses and bureaucracies of all types. By their nature, they are not a standard, though they are composed by use of standards. They are a way to use technology to seize opportunity more promptly. If you would like to take advantage of the experience of a company with deep experience with all the challenges you will face in embracing your service-oriented future, IBM is already there. The money you spend with IBM[11] may be offset handsomely by the pervasiveness of IBM's expertise and the way that expertise allows you to innovate without boundaries. **Consider your dreams for your enterprise. If SOA and infrastructure virtualization are needed to make them real, IBM can help you get there sooner, rather than later.**

[9] For more about IBM's Usage and Accounting Manager, see **The Clipper Group Navigator** dated July 3, 2006, entitled *Accounting as the Basis for Civilized Computing – IBM's Usage and Accounting Manager*, and available at http://www.clipper.com/research/TCG2006053.pdf.

[10] IBM's open sourcing of Eclipse and other intellectual property, its power.org collaborations, and its thousands of ISV partners are evidence of its collaborative achievements, and its ability to act as broker for the capabilities of others.

[11] Although IBM has some no-cost Community Edition products, they are a small part of its portfolio.

About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

Anne MacFarland **is Director of Data Strategies and Information Solutions for The Clipper Group.** Ms. MacFarland specializes in strategic business solutions offered by enterprise systems, software, and storage vendors, in trends in enterprise systems and networks, and in explaining these trends and the underlying technologies in simple business terms. She joined The Clipper Group after a long career in library systems, business archives, consulting, research, and freelance writing. Ms. MacFarland earned a Bachelor of Arts degree from Cornell University, where she was a College Scholar, and a Masters of Library Science from Southern Connecticut State University.

- ***Reach Anne MacFarland via e-mail at Anne.MacFarland@clipper.com or at 781-235-0085 Ext. 128. (Please dial "128" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, **The Clipper Group Explorer**, **The Clipper Group Observer**, **The Clipper Group** *Captain's Log*, **The Clipper Group Voyager**, and *"clipper.com"* are trademarks of The Clipper Group, Inc., and the clipper ship drawings, "*Navigating Information Technology Horizons*", and "*teraproductivity*" are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.