



Multiplexing the Data Center with SOA and Grid

Analyst: Anne MacFarland

Management Summary

Business IT environments continue to struggle with change. There are many kinds of processes, and each needs to be exposed to its 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 or fewer and different resources at some times than at others. The urgency that drives data center use is a matter of business demand, which cannot always be scheduled.

These characteristics are not well matched with the way IT organizations traditionally deploy, integrate, and manage their assets to deliver applications. Low-asset utilization and licensing charge overages are painfully obvious evidence of this mismatch, but there is more. The bigger and more comprehensive an application is, the harder it is to integrate or amend safely. This is not a matter of bad practices but of the heritage of focusing on the most expensive assets – servers and storage.

These days, the most expensive asset is time. Focusing on the time dimension of functionality (the service) rather than its physical dimension (servers and storage) exposes a set of metrics that allows the data center to align operations more closely with business requirements. Think of the changes that on-line tracking information has made to the shipping industry. Adding bar codes to packages and positioning devices to document their location enabled the shipping industry to enter a new stage of profitability. Multiple qualities of service could be supported. Instead of *as soon as we can*, we have *next day*, *three days*, and *economy* options – with rates to match. Customer expectations and satisfaction grew. One often-ignored benefit is the extent to which tracking improved the processes for the shipper. Fleets of vehicles could be used more effectively. Routes could be optimized before the packages arrived at their final distribution point.

For IT applications, *Extensible Markup Language (XML)*¹ is like the shipper's bar code. With it, 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. *Grid architectures* use XML to communicate between resources and to manage them. **Service Oriented Architecture**

IN THIS ISSUE

| | |
|------------------------------|---|
| ➤ The Focus on Service | 2 |
| ➤ Conclusion | 3 |

¹ Like HTTP, it is a common, standard information format. It supports self-description of unstructured and structured data. This self-description allows the data to be used by different information systems and to be shared across Intranets and the Internet. It supports any Unicode character and thus many alphabets. It is used in hierarchical schemas that can be understood by humans. It is widely used. Its shortcomings are that it is verbose, compared with some other approaches, and somewhat limited in its support for some IT concepts and nuances.

(SOA) uses XML to stitch components together into comprehensive applications. In both cases, the move from *hard-coded integration* to a *looser 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 allows the Data Center to initiate a new service on existing infrastructure as a matter of negotiation rather than as a powered-down reconfiguration.

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

Grid enables 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 to focus on business issues.

SOA is a software cousin of Grid, not a clone. Its focus is on deconstructing large brittle applications into more-easily-evolved components that 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 handcrafting of traditional enterprise application integration (EAI).

Grid's benefits, like time-division multiplexing over a telephone wire, make more efficient use of assets and enable their swap in/swap out via 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. Let's think about the

concept of service that underlies both architectures.

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. The focus shifts from often magnificent but rapidly-aging technology monuments to a focus on the services created and the infrastructure that best supports 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 by describing 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. While Grid has been used to embody processor and/or data virtualization, **what is really needed for full data center efficiency is virtualization of all infrastructure components – including network elements and workloads. This is more appropriately named infrastructure virtualization.** With it, when SOA brings together application elements in a runtime, they can be deployed where they provide 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 ITIL² but also growing out of development and testing organizations of various sorts.³ Industry standards, such as XML, create the opportunity for additional 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 coordination possible through

² ITIL (Information Technology Infrastructure Library) is a set of best practices for data center operations.

³ 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>.

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.⁴ This prudent longer-term continuity is inherently 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 effectively he can add and integrate new features. ***However, all the elegance you get in code-to-code integration costs you 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.⁵

Common Capabilities Needed to Civilize Distributed Systems

The worlds of hardware and software are very different. SOA and the infrastructure virtualization of Grid are no exception. However, they do use many of the same tools and techniques including 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. 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 supplies 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 for both SOA and Grid.

Because they share a common approach to common problems, using both architectures benefits from their synergies. 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 that meets 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 RAS⁶ features of the hardware and the runtime controls of the application.***

Conclusion

Service-oriented datacenter practices, in the form of Grid and SOA architectures, are a long-term boon to businesses and bureaucracies of all types. By their nature, they can be used to run tasks in parallel to speed time-to-completion of key business processes. By their nature, they can be used to reduce inconsistencies of process and over-subscription to software licenses. They offer many ways to use technology to seize opportunities more promptly. Think how you can benefit from using them in your organization.



⁴ This old approach becomes increasingly spendthrift, as assets get more reliable.

⁵ 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.

⁶ Reliability, Availability, Scalability.

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, Clipper Notes, 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.