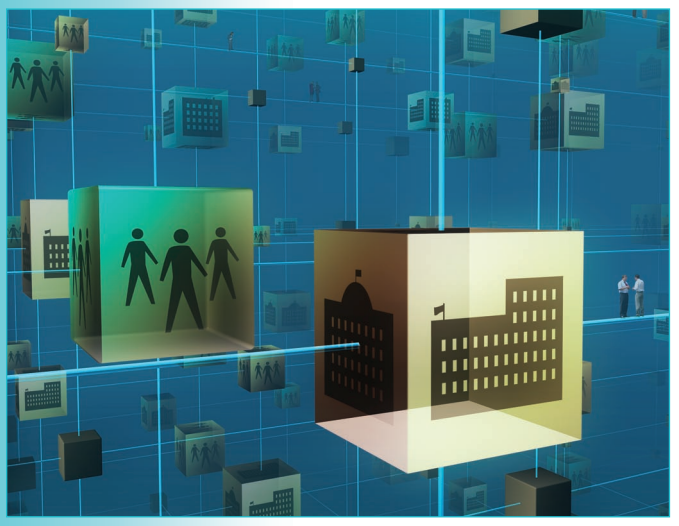

Virtual Organizations

Matei Ripeanu
University of British Columbia

Munindar P. Singh
North Carolina State University

Sudharshan S. Vazhkudai
Oak Ridge National Laboratory

Today's organizations are no longer constrained by traditional time and place barriers. Instead, information technology supports *virtual organizations*: flexible networks of independent, globally distributed entities (individuals or institutions) that share knowledge and resources and work toward a common goal. Generally, the set of shared resources isn't limited to computing power but includes elements as diverse as storage, network links, data sets, analysis tools, sensors, and scientific instruments. The sharing policies that govern the use of these resources are necessarily highly diverse, given that sharing must be flexible, secure, and controlled over time and place, while respecting users and purposes.

The Landscape

A number of research communities — preeminently grid computing and business-management communities — have embraced the challenge of exploring and creating the necessary infrastructure to support virtual organizations. They've made significant progress over the past decade: the virtual organization concept has been successfully realized in a growing set of domains ranging from collaborative work in business and science, to aircraft engineering, to multiplayer games, to medical data management, to supply networks. The software infrastructure enabling these deployments has become increasingly sophisticated, reliable, and scalable, and the research community is now better positioned to capture real users' requirements. At the same time, the attention has shifted and goals have become more ambitious: the early focus on providing a software toolkit to support the creation and operation of virtual organizations has moved to building a uniform infrastructure that, in addition to standalone software components, includes services and resources (computational resources, intermediate data storage facilities, high-speed networks, and so on). A telling example of this shift is the massive set of computational resources connected by optical links that the TeraGrid (www.teragrid.org) deploys to support several virtual organizations in collaborative science.

1089-7801/08/$25.00 © 2008 IEEE

Yet, the numerous challenges that have stymied wider adoption include the heterogeneity of existing resource-sharing protocols, the complexity of policies for sharing, the limited availability of enforcement mechanisms, the incomplete support for predictable quality of service and dynamic membership, and the social issues that relate to contract negotiation or risk analysis inherent to any large-scale deployment. The articles included in this special issue address some of these problems.

In this Issue

In "Bridging the Gap between Legal and Technical Contracts," Alvaro Arenas and his colleagues discuss the gap between the legal contracts binding participating institutions and end users and the technical agreements that can be established between the virtual entities that act on their behalf. The authors argue that legal contracts must translate into binding agreements at the virtual organization level. Supporting these agreements presupposes monitoring and enforcement mechanisms, and the virtual entities they link need to be traceable to real-world, accountable participants. The authors also discuss two aspects technologists often neglect: risk analysis for contract drafting and support throughout the virtual organization life cycle, from the identification of objectives and participants to the termination of the virtual organization.

A factor heavily impacting virtual organizations' current deployment is their integration with existing tools and operating environments. Existing grid toolkits generally follow a metacomputer design and are built as an integration layer that aggregates largely unmodified, operational systems. This approach often involves complex usage semantics and poses difficult integration questions with respect to bridging heterogeneous security and access-control mechanisms, auditing, and reconciling between local and virtual organization-level resource allocation policies.

In "Virtual Organization Support within a Grid-Wide Operating System," Massimo Coppola and his colleagues break with the metacomputer model and study the minimal functionality required at the operating system level to support resource aggregation. Key to the proposed approach is a layered system design and the reuse of authentication mechanisms available with existing operating systems for virtual organization support. In the long term, the debate that the authors initiate could lead to community-wide agreement on the set of primitive services that resources participating in a virtual organization should natively offer. This step is essential for simplifying virtual organization deployment and operations, thus widening adoption, and, in effect, is a natural step in networked systems evolution: once a function becomes widely adopted and its interfaces stabilize, it's either incorporated in the the middleware layer or outsourced to an infrastructure utility.

Providing adequate and predictable quality of service is a challenging problem for all networked services, not just virtual organizations. Given services' complexity and their dynamic nature afforded by on-the-fly service composition, even planning an adequate initial service deployment is difficult. Today's services often suffer from an inadequate ability to react and, once in operation, a burst of load can quickly lead to unacceptable performance degradation. In "Provisioning for Dynamic Instantiation of Community Services," Li Qi and his colleagues address exactly these challenges in the technological context of the Web Services Resource Framework and the Globus Toolkit. They present an architecture and preliminary experiments that demonstrate the ability to dynamically provision resources at various levels of aggregation, from single-node services to virtual organizations spanning multiple sites.

Despite remarkable progress, significant human and hardware resources must be allocated to support each virtual organization deployment.

These three articles represent only a sample from the broad range of ongoing research related to virtual organizations. Despite remarkable progress, significant human and hardware resources must be allocated to support each virtual organization deployment. The large numbers of functionally independent services that must

be deployed, configured, and reliably maintained require complex expertise and significant resources. As our understanding of usage scenarios matures and as standardization efforts, such as those led by the Global Grid Forum, gather traction, a uniform infrastructure that supports virtual organizations will emerge the same way today's utilities have emerged from a host of disparate and often competing efforts.

What we find most compelling about the virtual organization concept is that it shines a spotlight on the people-oriented — and in a broad sense, business-like — interactions that occur in open environments, whether the explicit goal be a traditional business deal in which each party seeks financial benefit or resource sharing among autonomous yet cooperating stakeholders seeking only to advance science. We believe that, increasingly, virtual organizations will be the way we understand, model, build, and administer large-scale computing systems.

Matei Ripeanu is an assistant professor in electrical and computer engineering at the University of British Columbia. He is broadly interested in distributed systems with a focus on self-organization and decentralized control in large-scale grid and peer-to-peer systems. Ripeanu has a PhD in computer science from the University of Chicago. Contact him at matei@ece.ubc.ca.

Munindar P. Singh is a professor of computer science at North Carolina State University. His research interests include multiagent systems, service-oriented computing and, especially, policy-based governance of virtual organizations. Singh has a BTech in computer science and engineering from the Indian Institute of Technology, Delhi, and a PhD in computer science from the University of Texas at Austin. Contact him at singh@ncsu.edu.

Sudharshan S. Vazhkudai is a research scientist at Oak Ridge National Laboratory, a US Department of Energy Laboratory. His research interests include distributed storage systems, grid computing, and high-performance computing I/O subsystems. Vazhkudai has a PhD in computer science from the University of Mississippi. Contact him at vazhkudaiss@ornl.gov.