META Group

META Delta

26 February 2003 File: SMS 1145

# Hyperdistributing Management Technology

**Service Management Strategies**

Glenn O'Donnell

**As infrastructure components become more adaptable and abstract, management technology is becoming increasingly distributed to instrument the changing dynamics of monitoring and configuration management. Using more localized data collection and processing nodes in cooperation with other management nodes is the trend for management systems. This hyperdistribution of management intelligence represents a radically expanded management architecture over current centralized systems. It is in its early stages and will not mature until 2008/09, but nascent segments of hyperdistributed management are emerging in network infrastructure and grid computing efforts.**

Current infrastructure and application management solutions conform to a centralized architecture, but distribution of selected management technologies (e.g., availability determination, performance analysis) is gaining popularity. This distribution trend will continue as infrastructure evolves through adaptive trends (e.g., grid computing), progressing through 2008. Under these changing conditions, existing management architectures will be strained, and eventually collapse, because a centralized architecture addresses dynamic local effects less nimbly. Given that such infrastructure inherently requires advanced embedded intelligence (e.g., sense and reconfigure), this intelligence will be exploited for broader management systems. Ultimately, embedded management intelligence will become pervasive (i.e., hyperdistributed), and infrastructure will become self-adaptive and exhibit organic behaviors (see SMS Delta 1026).

Initial technology developments are surfacing in pockets of network (e.g., Cisco's Internet OSS Programmable Network, Inmon's sFlow) and server (e.g., IBM's autonomic computing, HP's Utility Data Center) intelligence. Although emerging vendors will pioneer many hyperdistributed concepts, hardware vendors will control the underlying technology and overall progress. Vendors will embed advanced management software as an inherent component of infrastructure — not only supplying the data, but also applying correlation, knowledgebase, and adaptive control at the component level. Innovation is occurring first at the infrastructure level. Hyperdistributed application management will not begin until 2004/05, and it will require advanced application instrumentation technologies (e.g., JMX, .Net, Web services). Highly adaptive hyperdistribution with pervasive local intelligence will emerge during 2009-11.

A first step, automated infrastructure provisioning, is now a reality in some infrastructure pockets, though few IT organizations (ITOs — 5%-10%) are currently employing such technologies. By late 2004, infrastructure virtualization will emerge, with broad use by ITOs in 2005/06. During 2007/08, common distributed logic will span technology silos (e.g., server, network, database) via integrated policies and yield end-to-end adaptable IT services. Technology that enables infrastructure components to interact automatically with each other and with applications in cooperative automation will develop slowly, with organic infrastructure in early use by 2009.

Competitive differentiation will fuel infrastructure vendor innovation (with substantial software investments to create these hyperdistributed environments), but the aggressive nature of software evolution will render differentiation transient. Once a concept is proven successful, other vendors will rapidly follow, especially in the server domain. The network domain

***META Trend: During 2003/04, influence from non-IT personnel on infrastructure and application management (IAM) purchases and priorities (e.g., business views, user response time, service-level management) will accelerate. Through 2004/05, expanded instrumentation (e.g., Web services, Java Management Extensions, infrastructure vendor supplied) will continue to drive commoditization, leading vendors to explore automation of more complex processes (e.g., change, configuration) for varied infrastructure (e.g., desktop, networks, servers, wireless devices). By 2005, desire for data-level integration will spur added demand for XML as the management data exchange layer, but true mass standards adoption will not occur prior to 2007.***

will be less dynamic because network device processing limitations differ substantially across vendors and models. High-end network components (e.g., Cisco, Foundry, Juniper) provide adequate resources to support localized management. Each has unique capabilities based on proprietary technology, which will continue to provide leading-edge capabilities. Standards will lag proprietary options by two to four years.

Management agents, in their current form, will diminish as standalone products and gradually be replaced by embedded instrumentation (see SMS Delta 1143). Such instrumentation is beginning to incorporate capabilities that enable it to be remotely utilized in a secure manner and to apply local logic to the instrumentation data, which is changing the need for monitoring agent technology. The implication for distributed management is that embedded instrumentation increasingly executes more management tasks directly at the points of interest. Because management processing is tightly localized, many of the limitations of conventional data collection are overcome. Current systems rely heavily on polling (e.g., ICMP Echo, SNMP) to determine the state of the managed component. Polling is only as effective as the polling interval enables it to be. Common intervals are 10-20 minutes, which means visibility into availability and performance can be missing for up to 10-20 minutes.

As ITOs strive to measure more stringent service levels, such intervals become unacceptable. The proverbial 99.999% availability goal allows only 29 seconds of downtime per month. A polling-based system must use an interval of fewer than 15 seconds to be statistically accurate at detecting a single 29-second outage. Accumulations of shorter outages are even more complicated. This is unrealistic, because the demands on the polling station and the network path are excessive. The solution is to eliminate the polling by collecting and processing at the source, where measurements can be instantaneous.

A hyperdistributed management system requires enhanced resources (e.g., CPU, memory, cooperative communications) in the distributed components. However, Moore's Law is improving this situation, because silicon and architectural evolution will seemingly continue indefinitely. The overall hyperdistributed system is more efficient because unnecessary communications and processing at a central management station are minimized. Communication between hyperdistributed components and a central manager only transpires when data exchange is necessary to provide summary data or event notification. Current management architectures are inefficient consumers of data, with more than 90% of management data never used.

Reliance on central management processors will still be necessary, but a confederation of processors (e.g., midlevel managers) will gradually overtake monolithic servers. These processors can be relatively small, because their reliance on hyperdistributed intelligence will reduce collection and processing requirements relative to current architectures. Management appliances will gain popularity for such chores (see SMS Delta 1054), compelled by continued commoditization and simplicity requirements.

With the projected growth of Web services, management will become yet another Web service. Many of the same protocols will be used, and service directories will store macro-level configuration information to identify distributed management instrumentation and higher-level management processors (e.g., appliances). Relationships will be identified and tracked to guide cooperative interactions among hyperdistributed management components.

The similarity between IT environments and natural organisms will eventually produce interesting and useful patterns from interactions among hyperdistributed intelligent components. The science of emergence and complexity theory is being applied in academic settings to study how large collections of autonomous entities can yield sophisticated and elegantly organized collective behaviors despite the absence of a central authoritative force (e.g., ant colonies, weather patterns). Achieving such naturally interacting IT infrastructure will be science fiction through 2010, but results of early developments will be useful by 2006. Even existing self-adaptive technologies (e.g., OSPF network routing) exhibit characteristics of emergent behaviors.

## *Bottom Line*

**Demand for adaptive services and lower operational expenses, coupled with ongoing technology developments, is driving management technologies to higher levels of distribution. By 2010, self-adaptive infrastructure will be commonplace and incorporate high degrees of automation.**

***Business Impact: Enhancing the business value of technology can be achieved by investing in automation.***