the Availability Digest

www.availabilitydigest.com

Megaplex Modeling: The Future of NonStop Demand

October 2010

In its earlier papers, Megaplex: An Odyssey of Innovation[1] and Roadmap to the Megaplex,[2] the Standish Group traced the history of the Tandem computer from its development to its current incarnation as an HP NonStop server. The papers described the innovations brought to computer technology by the Tandem systems with their fault-tolerant multiprocessor architecture, and Standish envisioned a new NonStop architecture that it coined the "Megaplex."

The Megaplex

In Megaplex: An Odyssey of Innovation, the Standish Group defined the Megaplex as "a fabric of resources that will provide for application services for the next 35 years." It suggested that NonStop technology will be the basis for the Megaplex. The Megaplex will be a fabric comprising a wide variety of processors from low-end to highly redundant systems. Organizations will be able to manage their resources by dynamically deploying services to match performance, availability, and budgetary requirements. In the future, organizations will do away with periodic releases of massive applications and will instead have a continuous development process of smaller units of work. These micro-applications will be managed through the Megaplex fabric.

In the Megaplex, all databases will be equal and transparent. Data will reside in any database, including Oracle, Sybase, NonStop SQL, SQL Server, and DB2. Data will be reorganized to reflect the operational needs of the function (OLTP, ODS, BI, etc.). This reorganization will be accomplished by the Megaplex with little if any intervention from operations staff.

In the Megaplex, there will be no specific NonStop hardware. The innovative functional technology brought about by Tandem will live on to provide availability, scalability, and security of all of the resources within the Megaplex. Only that part of an application that requires high availability will be run in a NonStop environment, leading to significant cost savings.

The Megaplex will inherit HP's fully automated management capability, which HP calls Business Technology Optimization (BTO). BTO will benefit from the capabilities of NonStop, making the Megaplex even richer in capability.

The Megaplex will be the perfect platform for cloud computing. NonStop technology will offer Megaplex consumers the ability to choose higher reliability and greater security for those applications that require these attributes.

---

[1] Megaplex, An Odyssey of Innovation, *Availability Digest*; December 2009.
http://www.availabilitydigest.com/public_articles/0412/megaplex.pdf

[2] Roadmap to the Megaplex, *Availability Digest*; July 2010.
http://www.availabilitydigest.com/public_articles/0507/megaplex_roadmap.pdf

The Megaplex and Application Modernization

Continuing with its vision of the Megaplex, the Standish Group in its paper Roadmap to the Megaplex introduced its vision of the Megaplex architecture. The Megaplex will be a collection of server blades acting together as a single system using multiple types of operating systems, databases, and other computer resources. The Megaplex will operate Linux, Microsoft Server, NonStop OS, OpenVMS, and various types of Unix, including NonStop OSS. The Megaplex will also integrate databases such as Oracle, Sybase, DB2, NonStop SQL, and SQL Server.

The secret for being able to take advantage of this flexible architecture for critical applications is to decompose applications into services using the SOA (service-oriented architecture) model. Critical services can then run on highly-available systems, and less critical services can run on commodity servers.

Standish envisions a six-step approach to modernizing stovepipe applications to run in a Megaplex environment:

1. Migrate proprietary databases to industry-standard SQL relational databases.
2. Move from green screens to modern graphical user interfaces (GUIs) and a web presence.
3. Change the applications to reflect the features and functions of system-oriented architectures.
4. Upgrade from traditional availability and disaster-recovery methods to continuous availability and disaster tolerance.
5. Implement proactive security measures to replace traditional passive security and protection systems.
6. Modernize operations by moving from a vertical structure to a horizontal structure using horizontal service views.

Standish submits that shifting critical applications from running purely on NonStop systems to running in the Megaplex shifts much of the workload from higher-priced processing to lower-priced processing without sacrificing availability. It estimates that this move can save in the order of 40% for safety-critical applications, 35% for mission-critical applications, 33% for business-critical applications, and 20% for task-critical applications.

The Economies of the Megaplex

In its third paper, Megaplex Modeling: The Future of NonStop Demand,[3] Standish homes in on a specific architecture for the Megaplex and compares its cost with that of more traditional approaches. Standish envisions systems of HP blades that can run any operating system. For the Megaplex, Standish focuses on NonStop systems and Linux systems running on common blades. This means that any blade in the Megaplex can run highly critical applications in a NonStop environment and can run everyday applications in a Linux environment.

Keep in mind that applications have been broken into independent services (per the SOA model) that can be called by any other services. The services comprising an application are then classified as “black” or “white.” Black services are those that will cause great harm to an organization if they are not available. Services that are white will have a minor impact.

Two operating environments are created, one for black services and one for white services. The black environment provides high availability – that is, it runs on NonStop servers in a

[3] Megaplex Modeling: The Future of NonStop Demand, *Standish Group white paper*; 2010. http://h20223.www2.hp.com/NonStopComputing/downloads/Megaplex-Modeling%20final-Final.pdf

multiprocessing fault-tolerant configuration. The white environment will take the Google approach to availability, which means lots of cheaper Linux blades that can be replaced when they break.

Thus, a Megaplex complex comprises perhaps hundreds or thousands of blades, some running NonStop and some running Linux. Though the hardware for all of the blades is the same, running in NonStop blades is more expensive than running in Linux blades because of the licensing costs of the two operating systems. Thus, black blades are more expensive than white blades.

However, blades are interchangeable. If the critical application load approaches the capacity of the black blades, some white blades can be commandeered for black-blade service. In this way, critical applications can easily handle peak periods. During these times, white services either run slower because they are more heavily loaded, or they are terminated for the duration of the critical peak load, depending upon their uses. Thus, the blades in the complex can be continually reconfigured to reflect the load of critical applications, thus maintaining their performance.

Standish then takes this one step further. Rather than configuring for peak periods with some excess capacity that is the usual approach, Standish proposes that a system be configured for its median use. The need for excess capacity is replaced by the capability to abort noncritical services when necessary. Usage pricing can now be based on the actual use of black and white services rather than on reserved capacity, whether it is being used or not. Thus, pricing based on preconfigured fixed configurations is replaced with capacity-on-demand and utility pricing (what Standish calls the NonStop Demand pricing option).

Standish provides several pricing comparisons of the Megaplex versus other approaches. The comparisons are based on an extensive and up-to-date database of cost and availability statistics that Standish continually collects from industry interviews. Standish uses as an example a payment-processing application. It assumes that 25% of the Megaplex workload is operating in black (NonStop) mode, and 75% is operating in white (Linux) mode. The costs include acquisition costs, operating costs, and costs of downtime.

Some of Standish's results follow:

Megaplex	$1,688,000
HP NonStop Blades	$2,270,000
Unix Server	$2,589,000
IBM z10	$2,654,000
Linux (Google style)	$1,455,000
NonStop Demand Megaplex	$1,574,000

Summary

Under the vision of the Standish Group, the future in computing lies in the Megaplex. The Megaplex comprises a great number of identical server blades that can run multiple operating systems – NonStop for services that require high availability and Linux for less critical services. The Megaplex is sized to accommodate the median load. If the load increases, less critical services give up their blades to critical services to maintain critical service performance. Capacity is on-demand, and pricing is based on actual usage.

Megaplex Modeling: The Future of NonStop Demand was written by Jim Johnson, Chairman of The Standish Group (www.standishgroup.com). The Standish Group provides research services focused on improving project success through its CHAOS services and on enhancing the value of IT investments via its TCO/ROI benchmarks.