

Home
Archives
About Us
Advertise
Feedback
Subscribe

Home > In A Netshell > **Full Story**

Bits 'n' pieces of the Networked world

Grid, P2P and Distributed Computing

Our other Sites

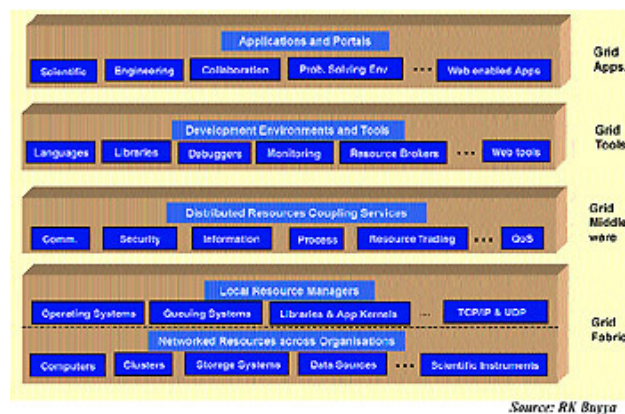


Figure 1.0 – Building blocks of a Grid computing model

It was quite sometime since I'd interacted with my friend and former colleague Rajkumar Buyya. He is currently in Australia, doing research at the Monash University. He called me during a visit to the United States, and updated me on his research.

His personal website and his resume stunned me, and the sheer variety and volume of the papers that he had published during the couple of years that we had lost contact simply amazed me.

His latest research pursuit is in the areas of Grid computing, which focuses on what they believe are fuelling areas of High Performance Computing (HPC) by utilizing a global network of computational devices, software and other instruments. His efforts in this direction also grabbed my interest. The biggest motivation for researchers of Grid computing is to address the drivers for HPC, which are notably areas like life sciences, digital imaging, CAD/CAM, e-commerce, aerospace and military applications. The Grid concept envisages the utilization of the distributed computing resources, to fulfill the needs of the larger community, by a global collaboration effort. The "Grid" in essence is therefore envisioned as an infrastructure that tightly integrates computations devices, software (perhaps in an ASP model) catalogued databases, specialized instruments, displays and people from widespread locations, and under different management authorities.

What is the kind of resources required to build a global collaborative mechanism? Programming tools, software that could translate requirements of an end application into the requirements of computers, network and storage, stringent security mechanisms, and end-nodes that can integrate tightly with heterogeneous, multi-speed, multi-protocol networks to start with. (See Figure 1.0)

There are major challenges in the actual application areas, concerning security (authentication, authorization) etc., resource allocation, scheduling, access of remote data sets or even collaboration on results that need to be addressed. I have not also seen a comprehensive report on how the bandwidth utilization and allocation takes place when disparate, distributed resources are being utilized, and this is something I'd like to look up in greater detail.

Research initiatives and test beds

However, there are dedicated teams working on these projects. Practical initiatives are being made in this direction. Perhaps one that has made headway is the Globus effort (<http://www.globus.org>), a joint initiative between the Argonne National lab, NASA, American national labs, and numerous universities. Essentially, the focus of Globus is on a large-scale computational grid, but side derivatives of this project are equally interesting. A key deliverable achieved by effort so far is the 'Globus toolkit' - a set of software tools that makes it easier to build computations grids, and power applications that run on them.

Validation of these efforts requires test beds. The National Partnership for Advanced Computational Infrastructure (NPACI) and the National Computational Science alliance have built large networked grids deploying the Globus software toolkit. A good source of information on Grid computing research efforts, and pointers to resources is available at <http://www.gridcomputing.com>.

Industry efforts

More interestingly, work on these areas has already moved from research labs, and test beds to commercialization. Grid computing, distributed computing, and P2P (peer-to-peer) computing - all sister areas, so to say-- form a triad of hot areas that are being focused on. The year 2001 is looked on as the year where distributed and P2P computing will come on its own. At least 15 companies focusing on distributed and P2P computing were funded in the last few months. Some call it "Internet computing". Others have catchy names of their own.

San Diego based Entropia (<http://www.entropia.com>) has developed Entropia 2000, a free software that activates and recycles the resources (CPU, Memory etc.) to quietly harness all the individual idle resources. (See Figure 2.0) Entropia has been in existence since 1997. However, there are upstarts who are jumping into the fray with gay abandon. Among these, most are companies looking at the software aspect, and not necessarily the Internetworking or bandwidth which is what I am concerned about. The general assumption is that the Internet is going to suffice. Recent entrants include Nextpage (www.nextpage.com), Distributed Science (www.distributedscience.com), Static Online (www.staticonline.com), and United Devices (www.uniteddevices.com). All these companies are leading efforts and contributing to the P2P effort. There are others in the stealth mode - Centrata (www.centrata.com) and Uprizer (www.uprizer.com) - who are yet to formally announce their products and business models, but are already drawing market attention. The prime motive is money, but a lot of big names and talent have been roped in into these companies from R&D institutions and Universities.

You could also volunteer to be a part of this global distributed computing network, either on a payment or non-profit basis. Process Tree (www.processtree.com) accepts applications from large organizations, as well as from individuals. Popular Power (www.popularpower.com) also follows a similar model to Process Tree, and offers to buy computing power. One other company that I must mention, and about which I will write sometime in detail later is Ejasent (www.ejasent.com), a company that has good management, investors and is aiming to gain a foothold in critical areas of Internet infrastructure, availability and performance.

Clearly both the industry, and the R&D community are sensing an opportunity here. Microsoft is endorsing P2P computing in a big way, Service companies like Viant (www.viant.com) are doing it, and so are vendor companies like Sun Microsystems. The latter is also contributing to Grid initiatives through its Grid engine software, and its 'Net Effect strategy'. Sun claims that its Grid engine software can effectively harness the resources on a network, and increase the efficiency by as much as 5 to 10 times the usable power of a network. It also claims to increase utilization levels to almost 98 per cent (<http://www.sun.com/gridware>). I am not too clear about the overheads involved at this stage, but as long as the efficiency claims are fulfilled, that is perhaps my prime consideration. Still, I would expect a large-scale deployment to involve considerable polling, to find out which resources are free, a scheduling mechanism that schedules and prioritizes jobs and an integration mechanism that functions in a non-proprietary environment. All these cause overheads on network utilization, and can cause performance bottlenecks. It will be interesting to see how these are resolved on a global scale.

It is ironic that Napster (www.napster.com), that helped create a new revolution to share MP3 files, and perhaps contributed to P2P computing getting a boost, has been sidelined. Improving upon this kind of technology are companies like Mojo Nation (www.mojonation.net) that are evolving innovative ways of distributing content.

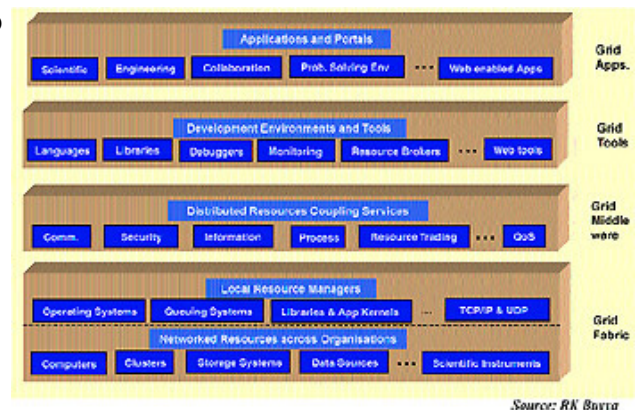


Figure 1.0 - Building blocks of a Grid computing model

Ultimately in the next one year or so, when these companies stabilize their solutions to address issues of distributed computing and content delivery mechanisms, they would need to synergize with players from the broadband communications space. With broadband becoming such a hot area, and broadband content delivery assuming considerable significance, this space is going to generate heat. Major broadband players are going in for the ability to offer enhanced personalization at the edges, and the end-user experience is going to be key. I strongly suspect that the two areas could converge sometime soon.

Broadband content delivery

The last two years have seen content delivery become a major force to reckon with. With the decline in e-businesses and dot.coms, this segment had been issued with a premature death warrant, but the industry suddenly seems to have come alive, and raring to go. Sometime back, I had dealt with content delivery (CD) network mechanisms in this column. I had also profiled some players in the CD market. Since then, there has been change in the way these companies are being perceived, not because they have lost the edge on technology, but because they have lost market momentum. While the market capitalization of companies like Akamai or Inktomi is nothing to write home about, with the steep falls in their net worth, the CDN community as a whole is very much active, with broadband equipment vendors, and internetworking giants like Nortel (www.nortelnetworks.com) and Cisco (<http://www.cisco.com>) emerging as strong players in this field.

The ability to deliver personalized content to the end-user, and provide premium services to generate revenue, however nominal, has been touted as the way to go, especially after all the free ISPs, and others providing free content, have been forced to close down shop, or reduce the scale of their offerings one after the other. The potential to generate revenue and then sustain it, is perhaps the only barometer of performance for a company, and rightly so.

But let me not digress. With the growth in the cable modem market, and end-users moving towards interactive television experiences, and personalized portal experiences, it is only right that broadband content delivery networks are assuming significance. The ability to personalize does not just mean customizing the look and feel, but also the means to provide a finer degree of access control (for the end-user and the ISP POPs), freedom to modify and access content with the right authority, and then to make all this accountable.

Consider the old Internet landscape (old in Internet time is an year ago) - See Figure 3. We had subscribers connecting to the Internet through ISP POPs over low-speed dial-up links, and accessing content hosts. There were numerous performance bottlenecks, right from the dial-up links which were of poor quality (more so in India), to the Internet backbone (a best-effort mechanism), to the inability of content-hosts to regulate and control the way in which content was delivered. Being transparent to the subscriber, his loyalty towards a particular ISP was motivated completely by his user experience. ISPs and portals took a big hit for sometime in their efforts to attract subscribers to their fold, and the cost of acquiring a new subscriber became quite expensive. ISPs, which began offering advertisement-powered free subscription as the ultimate incentive, have, as I mentioned been struggling to remain operational, and subscriber loyalty has been increasingly hard to maintain. The bluelights (www.bluelight.com) and NetZeros (www.netzero.com) of the world have been hard put. So is the case with the free PC, and free DSL providers. This has also adversely affected the large-scale hosting solution and service providers like Exodus (www.exodus.net) and Globix (www.globix.com), amongst others. In effect, it has been a chain reaction.

Now things are changing. The effort is no longer towards luring subscribers with the free carrot, but to offer service, and make the discerning subscriber pay for premium utilities and services. It is somewhat akin to what the cable TV medium has been following in the United States, with basic, medium, premium and personalized plans (the last is in the offing).

While it may take sometime for the web, and television mediums to merge cohesively, there are indications that we are going there. A similar approach on the Web needs to evolve and stabilize before the two can converge. How does one unify the end users, ISPs, hosting services, and content providers to deliver a consistent experience that benefits all? As a beginning, we see companies like Akamai (www.akamai.com), Sandpiper (www.sandpiper.com) and Digital Island (www.digitalisland.com) taking initiatives in moving the content closer to the user, from the content hosts to perhaps the localized ISPs, and their own POPs. This was essentially narrowband content, and the prime motive was to accelerate delivery of content, and enhance caching response times.

Broadband content delivery mechanisms are trying to build upon this. A range of features is


being added. This includes providing a granular degree of security, to controlled advertising, to quality of service, and pay-per-use mechanisms that could be dynamically activated.

How does one go about setting this model in place? What are the pre-requisites? For starters, subscribers need access to broadband connections, either through cable modems, fixed wireless access, DSL or whatever the mode might be. There is work going on in standardizing network logins, service advertisements, and service selection via a browser. Depending on the type of content, there could be mechanisms to deliver premium content over tunnels, to ensure a Quality of Service (QoS) look and feel. As far as the ISPs are concerned, or even the content hosts, work needs to be done on standardizing network logins (into the content site), and both the service advertisements and semantics need to be standardized to offer branded content services to the subscriber through a web-page. The ability to monitor and bill subscribers based on service level agreements (SLAs), and pre-defined policies is very vital.

The result is that broadband content is delivered to the end subscribers, directly via their individual broadband connections, and the ISP POP becomes a medium for content mediation. Broadband service nodes play a key role in these applications, and it is essential that both the service providers (viz. content, network and infrastructure services, and the vendors) viz. (Network, content caching and billing) join hands to provide the best possible end-user experience to the subscriber. I am happy to say that this is already happening.

Next month, there is a convergence of players in the content delivery space at New York (<http://events.stardust.com/cdn/>) and this should raise some dust. Industry leaders are participating, and hopefully, unlike the last couple of years, we should have technology, and not hype as the catalyst. There is an interesting, and extremely well written white paper on the website, which I'd recommend to enthusiasts of content delivery mechanisms. Look it up at http://events.stardust.com/cdn/documents/CDN_whitepaper.PDF
Till next time, Happy Networking!! NM

N. Shashi Kiran works for Nortel Networks at Santa Clara, as a Product Manager. The views expressed are his own. He can be reached at shashikiran_n@hotmail.com



[<Back to Top>](#)



© Copyright 2001: Indian Express Group (Mumbai, India). All rights reserved throughout the world. This entire site is compiled in Mumbai by The Business Publications Division of the Indian Express Group of Newspapers. Site managed by BPD