



Real-Time Infrastructure

*A Resource to Understand How Application Virtualization
Aligns IT with Business*

In this Issue

Introduction	2
Excerpts: "Extreme Transaction Processing: Technologies to Watch"	3
Improving High Availability through Application Virtualization	8
Real-Time: Q & A with LexisNexis	9
DataSynapse GridServer 5 First to Market with Mega-Grid Capabilities	12
UniCredit Case Study	13
Application Virtualization Improves Business Intelligence	14
DataSynapse to Host Inaugural User Conference	15
About Us	16

Executive Summary

Many organizations are turning to Application Virtualization to address business challenges that in turn will impact IT operations. What is application virtualization? How can it help my organization align IT with changing business priorities? Today's enterprises are always trying to stay one step ahead of their competitors and any type of program, promotion or are pushing to close more deals or reduce costs – all of these initiatives inevitably place a burden on the data center and the applications that support them. Application virtualization provides organizations the needed competitive edge. Business demand will always determine the allocated resources, and thriving organizations will always find a way to quench the demand. This newsletter will feature Gartner's Extreme Transaction Technologies report and explore how a Real-Time Infrastructure (RTI) and Application Virtualization impact your organization.

A DataSynapse publication featuring research from Gartner.

Introduction



Many organizations are turning to application virtualization to address business challenges that in turn will impact data center operations. What is application virtualization? How can it help my IT organization? I would like to share with you a short, real-time scenario that demonstrates how to align IT with changing business priorities through application virtualization.

To effectively compete and differentiate themselves in the market, a telecommunications company – Innovative Telco – makes a business decision to offer a new ground breaking program they are calling “Free 90-Day Trial – No Strings Attached”. The premise of the program is to get subscribers to switch from their existing phone service for a “Free 90-Day Trial” of Innovative Telco’s service and support. If the customer doesn’t find their service and support to be any better than the previous carrier, the customer has the option to switch back at no cost. Innovative Telco will pay the switching charges and usage for that 90-day trial period. This seems like a great deal – for the customer! The catch for the IT team: this program has to be launched in 45 days. Does this sound familiar?

At Innovative Telco, the customer-facing applications are static, siloed and running on a dedicated data center infrastructure. Innovative Telco is going to need to make significant changes to the application in an effort to handle the increase in application demand this program will bring and keep customers satisfied – past the 90-day trial period.

As a result, Innovative Telco internally raises several questions:

- How will IT support the volume of potential customers testing this service?
- Will the application supporting this program need to change?
- How will IT address proposed changes to meet the program delivery date?
- Will IT be able to make application changes in real-time, as demand increases, without manual intervention?
- Will IT be able to maintain the application’s need to be highly available, meeting the 99.99999% service levels or will IT need to take the application down to make the changes?

Innovative Telco is not alone. Their challenges aren’t unlike challenges other organizations face each and every day. Organizations are always trying to stay one step ahead of their competitors and any type of program, promotion or push to close more deals or reduce costs will inevitably place a burden on data center operations and the applications that support them. Application virtualization provides organizations the needed competitive edge. Business demand will always determine the allocated resources, and thriving organizations will always find a way to quench the demand.

The characteristics that define today’s successful, market-leading organization are:

- Innovative thinking, willingness to be aggressive
- Ability to differentiate from the competition
- Ability to adapt to changes in real-time, without delay or manual intervention

If this real-world example resonated with you, we hope that you will continue to explore the pages of this quarter’s issue, which focus on business initiatives like high availability, support for business intelligence and extreme transaction processing applications, in addition to customer case studies.

As always, we welcome your feedback.

Regards,
Jamie Bernardin
Founder and Chief Technology Officer

Excerpts from

Extreme Transaction Processing: Technologies to Watch

Extreme transaction processing (XTP) is an emerging application style aimed at enabling the implementation of large-scale, business-critical, transactional applications on the basis of distributed architectures implemented by leveraging commodity hardware and standards-based software. This research discusses a set of emerging XTP-enabling technologies.

CIOs, IT managers, architects, application managers, development managers and IT operation managers can use this research to learn how these innovative technologies can help them address the most-demanding performance, availability and scalability requirements, while keeping costs at a reasonable level.

Key Findings

- Increasingly, specialized startup companies, established platform middleware players and software megavendors are focusing investments on innovative XTP-enabling technologies. This is because they realize that previous-generation products (for example, transaction processing monitors [TPMs]) no longer meet users' total cost of ownership (TCO) expectations and that the current generation of Java Platform, Enterprise Edition (Java EE) or .NET-based products cannot support the most-demanding requirements.
- Incremental XTP technologies, such as computing appliances, distributed caching platforms and dynamic resource brokers, extend current products (such as enterprise application servers) with the goal of improving applications performance, availability or scalability.
- Disruptive XTP technologies – such as event-driven application platforms and grid-based application platforms – aim to go above and beyond mainstream platforms by proposing innovative programming models that extend the current ones with XTP-specific features.
- Most of these innovative XTP-enabling technologies are still in their infancy, have limited market proof and, in some cases, come from vendors of questionable availability.

Recommendations

- Mainstream organizations should primarily look at combinations of established platform middleware and incremental XTP-enabling technologies to address their emerging transaction processing challenges while minimizing technology risks.
- Leading-edge enterprises looking for sustainable competitive advantage should also consider the most-disruptive platforms. Even if their adoption is riskier, they enable the implementation of systems that cannot be supported at a reasonable cost on traditional, mainstream technologies.

STRATEGIC PLANNING ASSUMPTION(S)

By 2011, a new generation of application platforms stemming from the convergence of diverse XTP-enabling technologies will supersede Java EE and .NET as the platforms of choice for large-scale, business-critical applications (0.8 probability).

By 2010, support for advanced, event-centric programming models will be a standard feature in all application platforms aimed at XTP applications (0.8 probability).

ANALYSIS

1.0 What You Need to Know

XTP is an emerging application style addressing the modern breed of global-class, multichannel-enabled, highly distributed and highly interconnected, ultra-high-end transactional applications. Established transaction processing (TP) technologies such as TPMs or enterprise application servers (EASs) will be complemented and replaced in the long term by emerging technologies specifically designed to support the demanding requirements of modern XTP systems.

Organizations should prepare for this new generation of application by selecting the appropriate mix of incremental and disruptive technologies on the basis of their anticipated requirements and willingness to take risks to gain competitive advantage.

2.0 What Is XTP?

Gartner defines XTP as an application style aimed at supporting the design, development, deployment, management and maintenance of distributed TP applications characterized by exceptionally demanding performance, scalability, availability, security, manageability and dependability requirements.

Very much like traditional TP systems, XTP applications are aimed at enabling efficient, reliable concurrent and real-time access (read/update) to a shared database by executing application programs commonly referred to as "transactions." XTP applications are usually business-critical in nature. Therefore, they must provide transaction support as well as all the properties of traditional TP systems: availability, reliability, manageability, security, performance, scalability, accountability and extensibility.

In contrast to traditional TP systems – typically based on stove-piped, monolithic, and tightly controlled and "closed" architectures – XTP applications operate in a more complex and highly dynamic environment. They must support multichannel access, are based on componentized and microflow-enabled architectures, are designed to integrate with pre-existing applications in advanced service-oriented architecture (SOA) scenarios and are in many cases accessible by large, globally dispersed populations of often occasional users. Often, XTP applications also have stringent "time to deployment" requirements because they are meant to support fast-growing and fast-changing business strategies and emerging business opportunities. Moreover, XTP applications, unlike traditional online transaction processing systems, frequently deal with heterogeneous data sources (for example, relational data, Extensible Markup Language [XML] documents and event messages). Examples of XTP applications are e-commerce, Internet banking, online insurance, e-auctioning, online betting, electronic trading, online travel booking, microcommerce, micropayments, radio frequency identification (RFID)-based tracking and several other scenarios.

6.0 Dynamic Resource Brokers

Dynamic Resource Broker (DRB) technology – an integral component of Gartner's Real-Time Infrastructure vision (as a localized service governor operating at one specific software layer,

such as a DBMS, OS or application server) – is meant to help users optimize resource allocation and to further enhance scalability and availability for applications. One emerging form of DRB is specifically targeting support to mainstream EASs. Basically, a sophisticated form of loosely coupled clustering, this technology enables dynamic deployment or initiation of multiple instances of the same and other applications over a shared pool of networked computing resources, on the basis of defined policies. For example, to keep response time below a certain threshold for a certain mix of applications, at peak hours, more instances of a given application can be temporarily deployed on several servers (perhaps stopping or lowering priority and resource allocation of some other applications already running on those servers). When the spike of use is gone, the additional instances are undeployed so that the associated resources can be allocated to other applications that are, perhaps, dangerously close to reaching their critical performance or availability thresholds.

DRB technology can also be used to increase application availability, although its primary goal is to help users improve TCO by reducing “overprovisioning,” that is, the need to allocate – in an exclusive fashion – all the computing resources required to sustain the “worse case scenario” (as an example in terms of workload) for each application. Therefore, DRB technology provides a shared application infrastructure and resources, its allocation determined by policies.

By defining a shared pool of server resources to run in a loosely coupled cluster of shared computing resources, a DRB:

- Dynamically distributes requests to the appropriate server node (based on where the application instances are as well as their performance)
- Evaluates the performance of applications (typically through Java Management Extensions or proprietary instrumentation, along with defined rules and policies), spawning new instances when additional capacity is needed, or removing instances to make way for capacity required elsewhere
- Enables schedule-based capacity increases/decreases
- Handles resource failures by restarting application instances on alternative nodes, without manual intervention via policy-based mechanisms

This happens in a way that is virtually transparent to the applications, which therefore don’t need to be specifically designed to be deployed in a DRB-managed context. However, a modular and service-oriented application will better exploit a DRB environment than a monolithic application because it has a greater ability to be distributed and replicated.

DRBs are offered by specialized independent software vendors such as Cassatt and DataSynapse, although the larger EAS

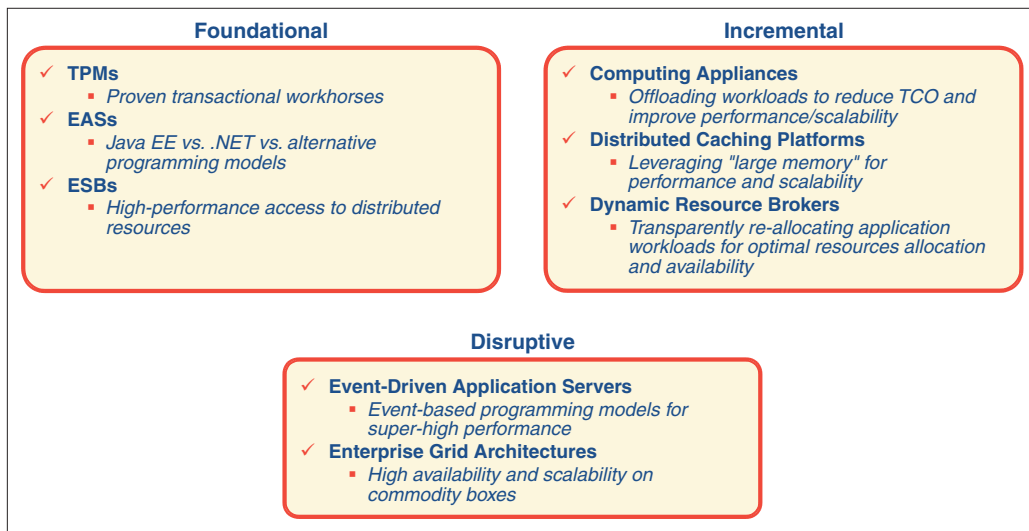


Figure 1. XTP Technologies to Watch

Source: Gartner (February 2007)

vendors (BEA Systems, IBM, Oracle and Fujitsu) are busy working to provide similar capabilities as part of their product portfolio (see Figure 4).

6.1 Potential Benefits

DRB technology enables users to improve availability and scalability of established EAS-based applications by optimizing the use of computing resources across multiple applications, thus alleviating the overprovisioning issue and ultimately reducing capital investments in hardware and software licenses. Depending on the effectiveness of management instrumentation and rules/policies, it could also decrease labor costs associated with manually rebalancing application resources and in dealing with resource failures (for example, installation/configuration/deployment of new servers). Because its use is virtually transparent from an application perspective, the introduction of DRB technology can happen once the application has been designed and deployed in production, with minimal or no disruption from the operational and software maintenance points of view. The DRB also has the potential to provide sensible data about the performance of business processes, which can be used by developers to refine and optimize applications.

6.2 Challenges to Mainstream Adoption

The current generation of DRB technology has some limitations:

- It can only manage the pool of resources it is given to manage because more resources cannot be dynamically added to the pool unless coordinated by a higher-level service governor.
- An application server-centric DRB manages computing resources (servers) only and cannot manage virtual networks or storage resources.
- It does not manage the end-to-end transaction environment (that is, the end-to-end response time) and just has knowledge, policies and optimization of the resources and section it knows about.

Initial implementations of DRBs will be largely based of homogeneous stacks from the same vendor (for example, BEA Systems, IBM and Oracle). Lack of technical maturity and a

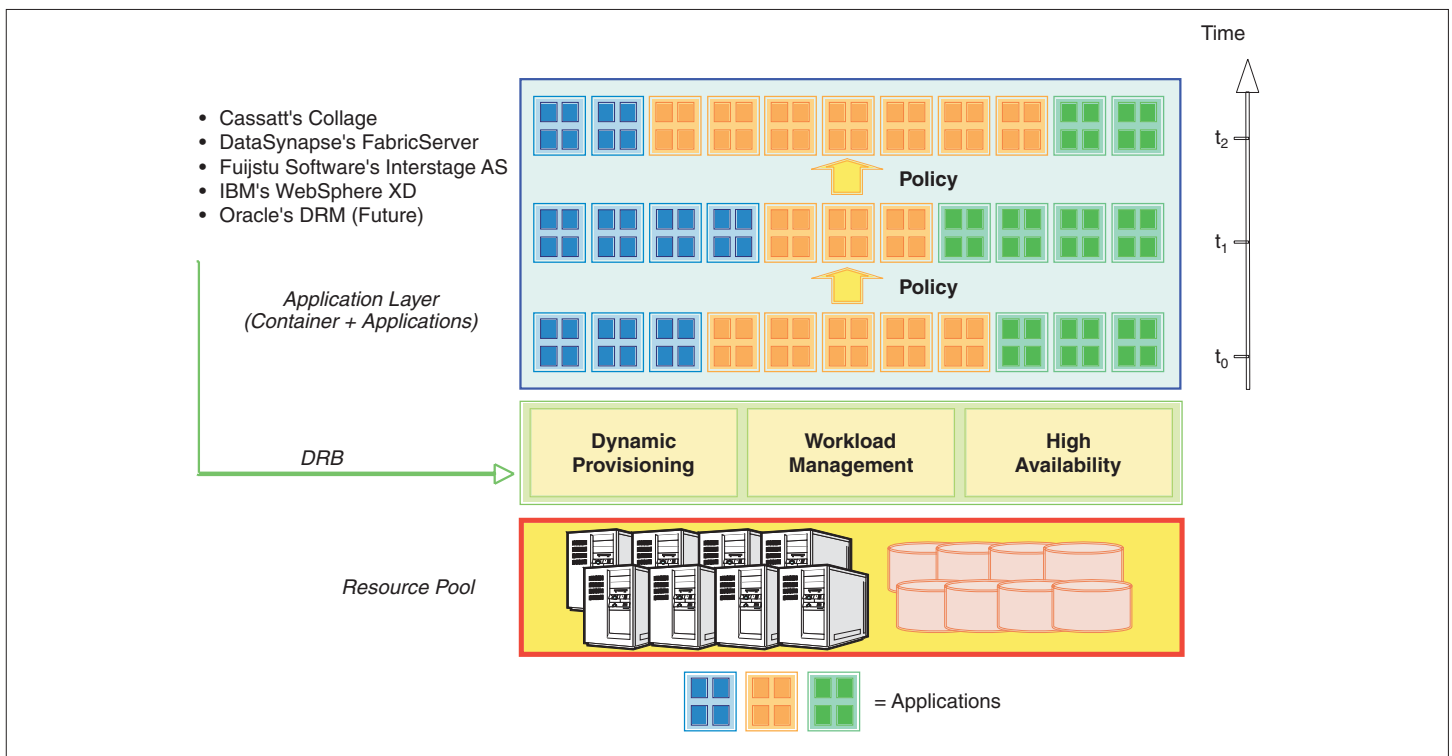


Figure 4. Dynamic Resource Brokers

Source: Gartner (February 2007)

more-evident and clear demonstration of TCO benefits are key obstacles for DRB's widespread adoption. Availability of specific powerful capacity planning aids; support for fine-grained security; detailed monitoring, management, logging and metering tools; and the emergence of commonly agreed-to standards for policy definition are requisites for market acceptance of this technology.

6.3 Action Item

Organizations seeking to reduce TCO and to improve scalability and availability of EAS-based applications should experiment with DRB technology, but also remember its still-immature state and the volatility of some vendors of these products.

8.0 Grid-Based Application Platforms

Grid computing refers to a collection of (typically commodity) servers from one or more companies that are coordinated to handle a common task. Grid computing has been used extensively during the past five years to create a large pool of computing resources leveraged to support batch-oriented, high-performance computing (HPC) applications, such as risk analysis in financial services, circuit simulation in the electronic industry or weather forecasting. However, the compelling price/performance of commodity Intel-based blades, high-speed network equipment and network storage devices is attracting the attention of companies interested in leveraging grid-like architectures to also support transaction processing workloads.

Grid-based application platforms provide an application runtime environment for service-oriented and event-based transactional applications, in addition to the same HPC, batch or analytical applications that typically take advantage of traditional grid computing. Grid-based application platforms' aim is to support "scale-out" scalability and maximum

availability on server-based "tera-architectures" (typically large clusters of Linux or Windows-based commodity servers). Typically, transactional applications in a grid-based scenario are designed as a collection of discrete, semiautonomous components glued together through a microflow language.

Grid-based application platforms leverage advanced workload management, peer-to-peer communication (such as Jini technology), distributed caching (such as JavaSpaces technology) and software deployment platforms (such as the Open Services Gateway Initiative Framework and Project Rio) to enable dynamic and transparent allocation across the grid (that is, on one or multiple servers) of as many instances of certain critical components of a transactional (but also HPC) application as needed to meet application-specific service-level agreements (such as response time, availability and so on). Also, grid-based application platforms typically enable hot plug-in and disposal of servers without stopping application operations, or manual reconfigurations by dynamically reallocating application components in an optimal way in the new configuration.

Most grid-based application platforms support a programming model that combines standard component software – such as Microsoft .NET, Plain Old Java Object, Spring and, in the future, Enterprise JavaBeans (EJB) 3.0 – and microflow technology (such as Business Process Execution Language or Web Services Choreography Description Language). Vendors offering grid-based application platforms include Appistry, Aumega Networks, GigaSpaces, Majitek and Paremus.

8.1 Potential Benefits

Grid-based application platforms can dramatically change the economics of even the most-demanding XTP applications by providing mainframe-like availability and scalability using

commodity hardware platforms. Their adoption will enable mainstream organizations to address business opportunities that previously could only be exploited by relatively few, technically astute organizations that could afford to adopt expensive and mostly proprietary application platforms (such as mainframe-based TPMs).

8.2 Challenges to Mainstream Adoption

Grid-based application platforms were released into the market from 2003 to 2005 by several small startup companies of little viability. Today, most products are available only in some geographic regions, and although rapidly evolving, some are still in early versions. Several proofs of concept, pilots or development projects were started during the past 12 to 18 months, and some notable projects have reached production. Nevertheless, whereas vendors of HPC-oriented grid-platforms (for example, Data Synapse and Platform Computing) have accumulated large installed bases, grid-based application platforms have only enjoyed a few dozen production deployments.

Another factor limiting the widespread penetration of grid-based application platforms is that applications must be developed according to their peculiar and unusual programming style. Many products are adopting popular component models (such as .NET, Spring and EJB 3 in the future) and standard flow languages; this will likely increase the appeal of grid-based application platforms by reducing the perception of risk and giving users the opportunity to leverage readily available skills.

8.3 Action Item

Leading-edge enterprises looking for advanced platforms to support XTP requirements should consider grid-based application platforms when their major concern is scale-out scalability, and nearly 100% availability for transactional and analytical applications (as well as hybrid combinations of both styles). However, also remember that these products have limited proof points in terms of real-life deployments, vendors are of dubious viability and the TCO benefits vs. traditional application platform claimed by vendors still has to be validated. Therefore, adoption of grid-based application platforms should be primarily considered only for specific scenarios in which established platforms are manifestly inadequate.

9.0 User Survival Strategies

Life for organizations, depending on XTP applications for competitive advantage, has never been easy, but now it's more complicated than ever. The old certainties, paradigmatically represented by transaction processing monitors, are increasingly proving to be functionally inadequate and too

expensive to maintain. New XTP-oriented application platforms are promising lower cost, greater flexibility and more power than traditional TPM, but are still unproven (for example, grid-based or event-driven application platforms) or not fully suited to the most-demanding requirements (such as EASs).

Many users face the two-pronged dilemma of meeting business requirements/competitive differentiation and displaying a willingness to take risks: Should users adopt highly differentiating, but of dubious viability, technologies or settle on the proven products, giving up the opportunity of achieving sustainable competitive advantage through the adoption of innovative XTP technology? (See Figure 7.)

Recognizing that there is no "one size fits all" XTP platform is the first step for organizations to resolve the dilemma:

- Despite the evident and continuous improvements of Java EE and .NET-based platforms, the most popular traditional transactional platforms – including most TPMs and CORBA products (such as Iona Technologies' ORBIX) or proprietary environments such as SAP's ABAP – remain viable products. They enjoy large, loyal installed bases that will continue to leverage their proven technology to run established business-critical TP systems, as well as to deploy new applications. These traditional platforms constitute a large and profitable, albeit modestly growing, business.
- It is in the vendors' interest to continue to support, improve and extend their traditional platforms with additional capabilities and tools. For example, most TPMs support modern technologies, such as Java, XML and WSs that in many cases users will use to expose TPM applications as high-performance, highly available services that will be consumed by new XTP applications.
- The TPM vendors' primary goal is to slow down installed base erosion, not to get new accounts. TPM customers require improved quality of service and TCO, not extensive user-

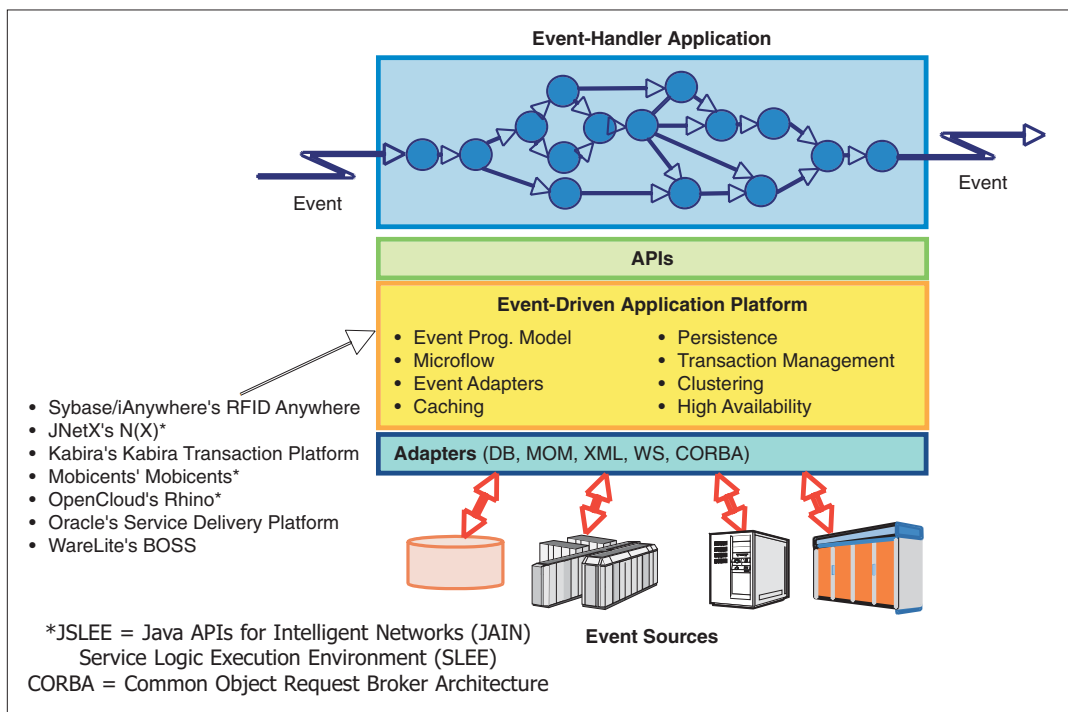


Figure 5. Event-Driven Application Platforms

Source: Gartner (February 2007)

Acronym Key and Glossary Terms

API	application programming interface	MOM	Message-oriented middleware
CORBA	Common Object Request Broker Architecture	ODBC	Open database connectivity
DBMS	database management system	RDBMS	relational database management system
DRB	dynamic resource broker	RFID	radio frequency identification
EAS	enterprise application servers	RPC	remote procedure call
EDA	event-driven architecture	SLEE	Service Logic Execution Environment
EDAP	event-driven application platforms	SOA	service-oriented architecture
EJB	Enterprise JavaBeans	SQL	Structured Query Language
ESB	enterprise service bus	TCO	total cost of ownership
HPC	high-performance computing	TP	transaction processing
Java EE	Java Platform, Enterprise Edition	TPM	transaction processing monitor
JAIN	Java APIs for Intelligent Networks	WS	Web services
JDBC	Java database connectivity	XML	Extensible Markup Language
JSLEE	JAIN SLEE	XTP	extreme transaction processing
JVM	Java Virtual Machine	zAAP	zSeries Application Assist Processor

experience features. Thus, TPMs will continue to enhance their dependability, scalability, performance, manageability and availability, but will not incorporate the rich, modern TP features available in more-modern platforms. However, they will remain the best option for risk-averse organizations and will participate in complex XTP scenarios by providing support to core, data-centric business services.

- EASs offer a good compromise between functionality and risk, but the most-sophisticated and demanding scenarios can be better supported through grid-based application platforms (when the main concern is availability and scalability) or through event-driven application platforms (when performance and low latency are the crux of the problem), although most mainstream organizations will consider these platforms too risky, given their immaturity.
- The combination of proven and advanced technologies is, in many cases, the most pragmatic way to take advantage of innovation by keeping risk at a tolerable level. Therefore, for mainstream users embarking on new transactional application

projects, an architecture that combines Java EE or .NET with a choice of incremental XTP technologies (distributed caching platforms, DRBs and computing appliances) will be the most reasonable and viable option through 2012.

XTP is a complex problem and cannot be solved with one specific product or technology. Traditional TPMs, EASs and the innovative, but immature, technologies described in this research will play a role and can be used in combination to solve real-life problems on a pragmatic basis. Most organizations will follow an incremental approach by complementing established and proven products with the appropriate combination of innovative technologies as needed for supporting emerging business requirements. However, in some cases they will have to adopt the riskier and more-disruptive technologies to grab greater business benefits and sustainable competitive advantage.

Excerpt from Gartner RAS Core Research Note G00146107, Massimo Pezzini, Milind Govekar, Yefim V. Natis, Donna Scott, 13 February 2007

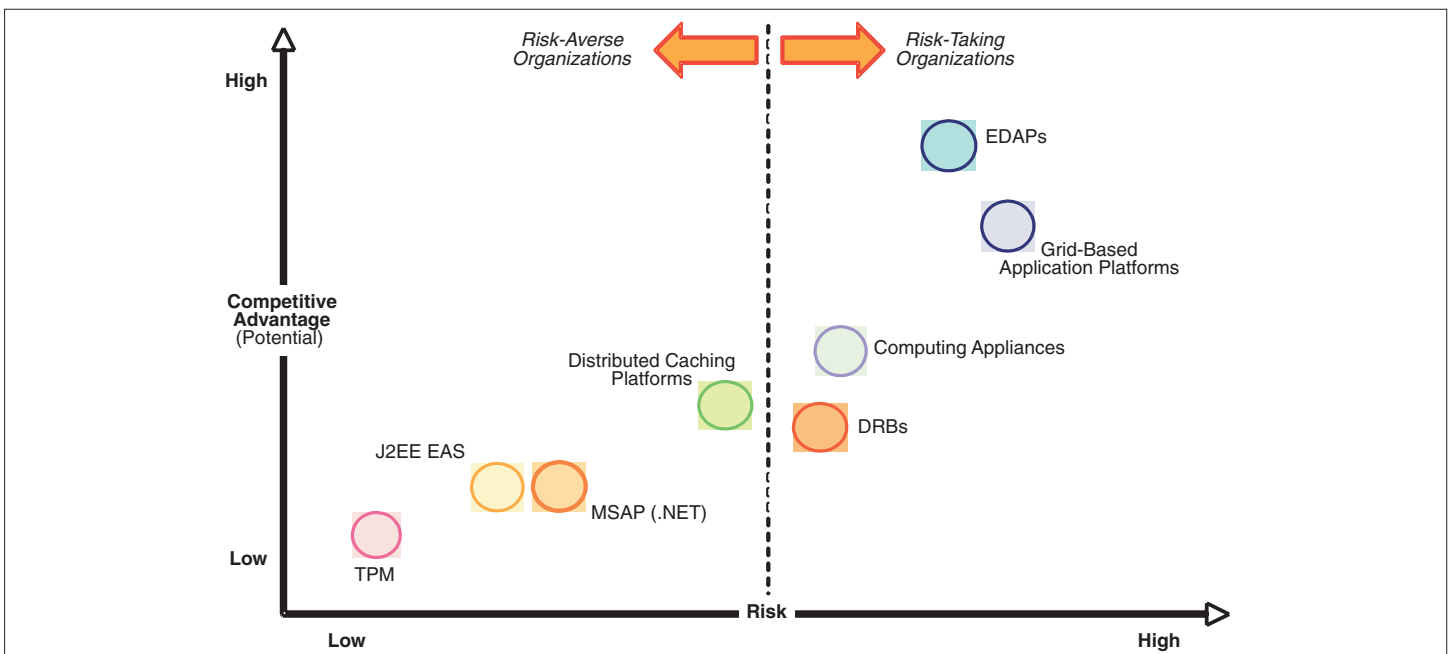


Figure 7. XTP Technologies: Risks and Rewards

Source: Gartner (February 2007)

Improving High Availability Initiatives through Application Virtualization

Business demand is dynamic – and often unpredictable. IT has to be one step ahead of the business, able to accommodate growth and shift capacity to different applications in real-time to meet the needs of customers and employees.

Given the integral role technology plays in daily operations, businesses have a choice: either accept the risk of application downtime or overspend to ensure that systems remain available. Both choices come at a price. Accepting risk could lead to slow or unavailable applications, usually during a spike in use. Essentially, businesses stand to lose the system when they need it most. Alternatively, they can over-provision, investing time and capital for systems and applications that will run most of the time at a fraction of capacity.

Because most firms are naturally risk-averse, larger IT investments are made to maximize up-time. High availability (HA) configurations seek to bring downtime to zero, as tolerance for service or application interruption becomes slimmer. As a result, businesses routinely over-provision hardware and software licenses in mixed-architecture environments to ensure that mission-critical systems are always available. Proliferation of these heterogeneous and redundant “silos” causes costs to escalate rapidly.

The Cost and Complexity of High Availability

HA environments usually consist of custom scripts developed over time to address specific problems. This reactive approach leads to excess technology, financial, and human capital investments. HA is costly, requiring extra servers and software licenses simply to implement the environment. This increases the financial outlay both up front and over the life of the equipment and software. With maintenance fees on software licenses creeping past 20% per year, the actual spend can be twice as high as the initial license costs.

The complexity of the HA configuration coupled with the priority of applications supported by HA require more employee time. System and application administrators spend more time managing the environment while the IT team spends time keeping existing systems running at the expense of new initiatives that can fuel company growth. The opportunity cost is real.

Because business demands it, up-time trumps efficiency. HA environments tend to rely on manual processes, static configurations and loosely documented custom scripts to ensure that an application will continue to operate when planned or unplanned changes occur. The cost is unnecessarily inflated, eroding margins and diverting skilled IT resources from value-creating endeavors to more tactical support roles.

The Virtualization Advantage

Fortunately, virtualization is coming to the forefront of advanced availability technologies and is reducing the inefficiencies and over provisioning of hardware associated with HA environments. Virtualization technology is enjoying a period of explosive growth, and increasing numbers of enterprises are becoming virtualization converts.

In HA environments, which require planned over-provisioning, applications are supported with excess hardware and are allocated additional software licenses. Provisioned for “just in case” scenarios, these resources often go unused or – at a minimum – are underutilized. The initial provisioning cost is exacerbated by the ongoing maintenance fees and internal IT management costs associated with having equipment deployed regardless of the extent to which it is used. Virtualization can obviate the need for planned over-provisioning with its nascent ability to pull resources from disparate platforms when needed and dynamically balance workloads across the environment. HA is maintained with fewer servers and at a substantially lower cost.

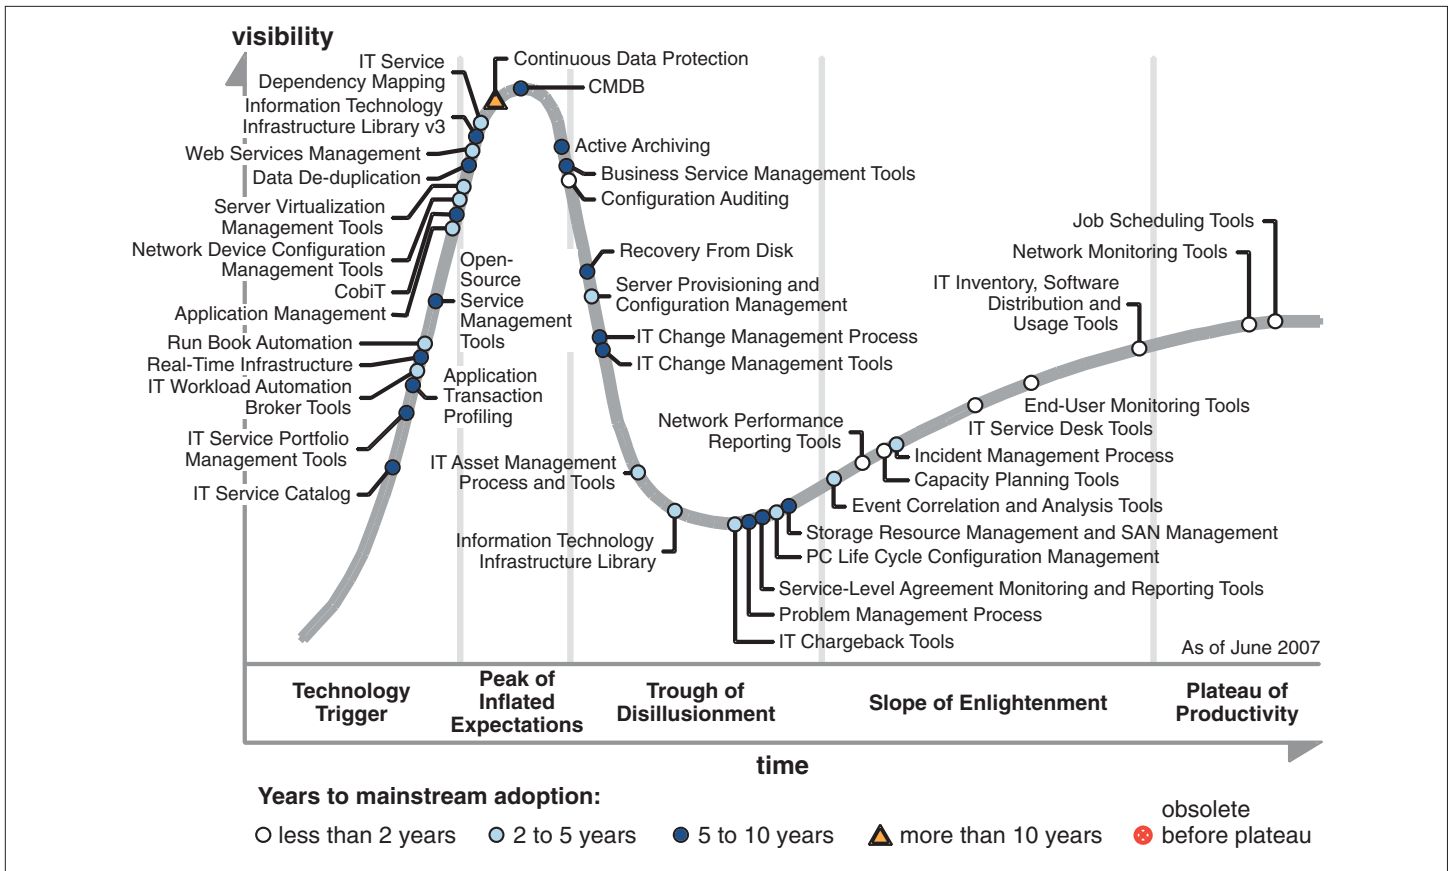
DataSynapse is a leading global provider of application virtualization software. Its FabricServer™ technology centralizes the command and control of application deployment and execution by virtualizing applications and guaranteeing that capacity is available on demand.

By expanding and contracting the use of infrastructure resources dynamically in a mixed-architecture environment, FabricServer delivers economies of scale through the more efficient use of available technology. FabricServer creates a real-time infrastructure by a pooling and adaptively allocating resources across the data centre, harnessing the power of the entire infrastructure to support specific applications. The flexibility afforded by FabricServer streamlines application and system utilization, reduces operating costs, and contributes to a HA environment without introducing additional complexity. With virtualization solutions, deploying extra hardware and purchasing additional software licenses becomes unnecessary.

For IT organizations, introducing a virtualization solution to automate and dynamically provision resources in a shared environment presents both a cost advantage and a significant ROI opportunity. Such a solution not only ensures HA capabilities for enterprise applications, it enables more efficient use of existing systems and application resources.

Source: DataSynapse

IT operations managers are challenged to "do more with less" (that is, deal with increased technology complexity, provide innovation and improve service under the constant pressure of reducing costs). Real-time infrastructure is "On the Rise".



Hype Cycle for IT Operations Management, 2007
 Excerpt from: Gartner RAS Core Research Note G00148631, 22 June 2007

Real-Time: Q & A with LexisNexis

Legal Data Meets the Grid

In this interview, Grant Peterson, vice president of product development for the LexisNexis Litigation Services group, discusses how the company's Applied Discovery application is able to handle twice the number of legal documents in the same timeframe since moving to a grid environment powered by DataSynapse and its application virtualization technology.

QUESTION: To begin, can you explain LexisNexis' Applied Discovery application?

GRANT PETERSON: Applied Discovery is an online service that takes in large quantities – I mean very large quantities, in the terabyte range – of document material, processes it, turns it around, and organizes it for review by legal teams prior to making discovery productions to opposing counsel in active legal matters.

The complexity of our business is that we are dealing with large numbers of electronic and scanned paper documents on short periods of time, needing to organize, index, sanitize, pull the metadata from and present for review that quantity of data.

Generally, the expectations are that we take in multiple hundreds of gigabytes, if not terabytes, of data and turn it around in days, so that it can be reviewed and then subsequently produced to opposing counsel. So, the volumes and turnaround times are our major challenges.

QUESTION: What was your system for handling this process prior to moving to the virtualized DataSynapse platform?

PETERSON: We have, effectively, three systems. Our initial system, the first point of technology application to the problem, is our collection platform. The data collection platform is used to gather the information that needs to be brought into a discovery matter.

We then have a production system, and our production system takes the collected material and essentially uploads it ... allows us to organize it into an intelligent review by what we call a "custodian," which is ultimately where the data came from – some individual, presumably. We upload that quantity of data and pull the metadata off each of those documents – and we

have many collection types, like Zip files, PSP mailbox files and LotusNotes NST files. We will pull those apart into individual messages, preserving the metadata on each of them, and store that metadata for what's called "non-spoliation." You need to be able to review the documents without changing the metadata and then deploy them for review in a searchable index.

Our solution prior to DataSynapse was based on taking a large number of stand-alone applications, or stand-alone instances of our production environment, and we would have a production engineer running maybe five copies of the production environment and monitoring each of them as they uploaded and preserved and indexed incoming files. So, the state-of-the-art for us prior to this was individual people, or human labor, monitoring multiple running applications. Somebody who was really good could perhaps monitor 10, but most people probably could look after five converters at a time. So, what we did was just have a large number of blades and a virtual terminal application, and a production engineer would essentially keep flipping back from one blade to the next, monitoring the conversion status on each blade.

QUESTION: How long have you been using DataSynapse GridServer?

PETERSON: We implemented GridServer last year and it is, at this point, fully deployed. What we did was take those individual production consoles and virtualized them into the grid and created a pipeline manager. So now, instead of a production engineer monitoring multiple jobs individually, what they're doing is monitoring a console of the full grid. We have about 350 blades that are active in our production environment, and we can now take about five production engineers and monitor the activity on the entire grid.

QUESTION: What were your biggest market challenges, and how did the GridServer platform help you overcome them?

PETERSON: For one, we've worked with our customers over the years to essentially come up with what we believe is the best review experience in the industry. That review experience is based on taking every document – all 140- some-odd disparate document types that tend to come in during discovery projects – and coming up with a uniform way to view, review, redact and deliver those documents. We standardized on converting documents to PDFs as a neutral review format. It allows us to preserve the original nature of the document across a very large number of initial document types and make the review experience efficient and consistent.

In order to do that, we have to convert huge numbers of documents to PDF, and there is competitive pressure on us to not add additional cost to our service as a function of making that additional conversion happen; some of our competitors review documents in text rather than in a more natural or native format. In order to convert the documents and not add

cost, we had to find an automated way to do it in large volumes that didn't drive up labor. So that's one factor.

The second factor is the volumes of data coming in during discovery are ballooning. They're ballooning, but the turnaround times expected and the cost of doing electronic discovery conversion are really staying a little more level. So, there's pressure to be more efficient and pressure to continue to be as fast as we were in the past. A year ago, we were handling about half as much data on a daily basis, now it's not unusual for us to turn around 6 or 7 million documents a day. It would have been nearly impossible given our prior labor model to do that without having hundreds of people monitoring converters.

By implementing GridServer, we've been able to keep our conversion labor consistent and double the throughput.

QUESTION: Have customers noticed your consistency even as document volumes keep rising?

PETERSON: We've always offered a best-in-industry service, but our challenge was expectations of price going down and the necessity, from a business growth perspective, of volumes going up. In order to maintain that excellence of service, something had to give. I think the feedback we've gotten is that we continue to be one of the fastest turnaround time vendors in the industry, and continue to meet customer expectations. Doing that at twice, or more than twice, the expected turnaround is validation in itself. I don't know that we've improved our service, per se, but we've doubled our service and maintained the same levels of excellence.

So, I think it's been transparent to our customers, but I believe that had we not made this shift, our customers would have noticed negative impacts on their expectations. What we've done has been to maintain levels of excellence by innovating in front of what are some difficult pressures in terms of turnaround times and data volume increases.

One other comment I would make is that I consider the production team – we have a large production team that handles the daily service deployment of content and customer projects – as much my customer as external customers, and we have gotten significantly positive feedback from them because the process of monitoring the pipeline against the virtual server is a significantly more accurate and simple procedure than was flipping from screen to screen to screen to locate the status of a running application. So, our internal customers are elated with the simplicity and, effectively, just the capacity that it makes it possible for them to stay in front of.

I also would say that the accuracy of the work they're doing has improved because in virtualizing this network of production systems, we've also done a much better job of managing errors and a much better job of recovering and restarting. In the past, if a

production engineer encountered an error in one of these consoles, that blade could be sitting idle for any period of time until they noticed and intervened and made whatever correction was necessary. In the new implementation, the system is self-correcting, and that blade gets put back to work immediately. We're actually getting a lot more capacity out of our hardware resources than we were in the past because we're able to keep the blades running.

QUESTION: Do you expect your current DataSynapse platform to scale as your demands in terms of data and turnaround time continue to increase?

PETERSON: We don't think we're anywhere near maxing out the solution. The first transformation was being able to efficiently deploy the people resources we had. The second transformation that we've encountered is being able to keep the deployed hardware at a higher level of load, and that's given us a decent amount of headroom in terms of this conversion process. In terms of next steps, to stretch this solution even farther, we've only taken two or three of the critical processes to the virtual grid at this point. We have a large effort underway called the "production control system," which is taking additional processes – including the import of data, export of data and production-set management – online and putting them also on the virtual system.

So, there's more in our production operation that we can virtualize, and, once we get through that, we can certainly add hardware to take up additional capacity requirements. One of the things about this solution that's wonderful is that it's quite scalable. At this point, I don't have any concern that this won't stretch to cover our business growth. That puts us in a very comfortable situation.

QUESTION: When will "production control system" be complete?

PETERSON: Project work is underway now; we're kind of in a rolling deployment. I expect we will be done with the full deployment of our entire production system on the virtual technology in early 2008. We have about 40 tools in our production environment, and we're dropping them at a rate of about two to three per month onto the virtual system.

QUESTION: Did you look at any other solutions before choosing DataSynapse?

PETERSON: We had four solutions in the arena when we made the evaluation. Unfortunately, I don't have the details on why we didn't select the others, but I know why we did select DataSynapse. Probably the primary reason [for selecting DataSynapse] was that it tested out to be the most mature technology that we evaluated, and it was more flexible in terms of allowing us to create a custom pipeline application than the other solutions that were evaluated. So, the programmability and the integrability in our core technology were the principle selection criteria, as well as just the track record and ... [the fact that] the relationship with DataSynapse was the best. There was a strong desire on both sides to work together based on what I'd consider to be good, interpersonal dynamics.

QUESTION: In a nutshell, can you summarize your overall experience in making the switch to a virtualized grid environment for your production system?

PETERSON: I'll start by saying – and this is unusual – that my expectations have been fully met, and exceeded, in this implementation. We made this implementation out of necessity, and the relationship with DataSynapse has been excellent from the start. The technology has performed as expected and as tested; we really haven't encountered any issues requiring escalation to my level to become resolved. I'm certain there probably were [bumps in the road] in doing the integration, as there always are, to get things working, but it got done on schedule; it got done without escalation; the relationship has remained strong; and there's really little question that there's excellent value in our investment at this point.

I consider the investment not just the cost of acquiring the DataSynapse technology, but also to include the cost that we laid in terms of redeveloping our system to run on the grid. In both cases, we've performed within budget, and it's enabled us to deal with 200 percent the amount of data we've dealt with in the past without having to grow our production organization to handle it. In that regard, I guess I would say it's been invaluable to us. I work with a number of vendors around licensed technology, and this one has been far and away the smoothest.

To read the full announcement, visit:
<http://www.datasynapse.com/LexisNexis-Implements-App-Virtualization.aspx>.

Source: DataSynapse

DataSynapse GridServer 5 First to Market with Mega-Grid Capabilities

Award-Winning Product Simplifies Enterprise-Wide Administration, Offers Enhanced Role-Based Security and Low Latency Capabilities

DataSynapse has announced the newest version of GridServer®, making it the first product on the market to offer mega-grid capabilities for grids of up to 20,000 nodes. GridServer 5 answers the call for improved global administration, including seamless integration with enterprise authorization frameworks and role-based security policies that allow organizations to protect application integrity in a distributed environment.

These major enhancements give organizations and administrators first-in-class management capabilities based on real-world customer experiences from the DataSynapse client base, many of whom are running global grids of 5,000+ CPUs that span multiple data centers.

“With scalability on demand, businesses can optimize the utilization of internal resources, reducing cost and complexity as well as achieving greater business agility,” said Mark Blowers, senior research analyst at Butler Group. “The end result is to make current applications meet service levels and the objectives of the organization.”

DataSynapse’s solutions help today’s enterprises achieve business-aligned IT by virtualizing, distributing and executing application workload in a guaranteed, scalable manner over existing computing resources. By matching application demand with supply of available system and data resources, GridServer provides an “on-demand” environment that dramatically improves application performance, resiliency and uptime. Organizations can dynamically scale any application service, at any time.



“More than a dozen DataSynapse clients are in production today with mega-grids that exceed 3,000 CPUs, with plans to grow to 10,000 and higher,” said Jamie Bernardin, CTO and co-founder of DataSynapse. “With GridServer 5.0, DataSynapse significantly simplifies the ease of administering a global, real-time infrastructure.”

Source: DataSynapse

Key Features of GridServer 5:

- “Low touch” grid administration capabilities that simplify management of mega-grids, including:
 - Enterprise Management Dashboard with Search – GridServer administrators can more easily track, visualize and search grid activity categorized by component, and can use grid-wide searches to quickly locate engines, daemons and services
 - AJAX-enabled dynamic views of current state grid information
- Advanced optimization of engine assignments with property-based affinities that enable preference-based workload scheduling
- User-Defined Roles for fine-grained security and entitlements with seamless integration to existing authentication and authorization systems
 - LDAP and Active Directory integration with group/role mappings
 - “Maker-Checker” functions to enable audited coordination and workflow of grid administration of key functions like provisioning and deployment of services
- A low latency .NET virtual service implementation is available with the GridServer 5.0 release that boasts a 2 millisecond task latency that can improve efficiency of fine-grained workloads

UniCredit Case Study

UniCredit Group is one of the largest banking and financial services organizations in Europe with a network of 7,200 branches and strong local roots in 20 countries. Its international network is made of branches, representative offices and small banking subsidiaries in 40 countries worldwide.

With the passing of International Accounting Standard (IAS) 39 and Basel II regulations, European banks are now required to establish principles for recognizing and measuring corporate holdings. These holdings include financial assets, liabilities and contracts to buy or sell non-financial items. The regulations also require that banks provide uniformity with how they approach risk management across nations. With these strict compliance deadlines, UniCredit's Capital Allocation Unit – Group Pricing department (CAGP) needed a system that could uniquely generate financial figures and perform computations in the least amount of time possible.

The central pricing application – named "Bulldozer"- had the goal not only to deliver timely and compliant financial figures, but also to reduce the firm's exposure to operational risks.

To deal with such a complex international compliance issue, UniCredit was prepared to face an ongoing, time-consuming project. The firm had to investigate best practices and cost-effective methods and evaluate vendors from the ground up.

The bank's IT department, UniCredit Global Information Services (UGIS), turned to DataSynapse, a global provider of application virtualization software, to achieve its compliance goals. DataSynapse was UGIS' first choice because of its international recognition as a market leader in financial services virtualization technologies. DataSynapse had also successfully helped other major international financial institutions with compliance challenges.

DataSynapse GridServer® is application virtualization software that enables application services to be distributed and executed across a shared, heterogeneous infrastructure. By creating a virtualized environment that matches user demand with supply of available system and data resources, GridServer provides an "on-demand" environment that dramatically improves application performance, resiliency and uptime.

"The enhanced computing power GridServer provides our pricing software has exceeded targets beyond our expectations. "Our original goal when we embarked on this project was to independently price 600,000 deals. We are now doing more than two million deals in the same timeframe."

Alberto Ricciotti
Head of Group Pricing,
Capital Allocation Unit
UniCredit

UniCredit developed its "Bulldozer" pricing application with GridServer to optimize processing of its intraday and end-of-day financial results. Bulldozer is a custom-built Java application developed to compute a high volume of financial figures for risk analysis and IAS 39 compliance. The processed figures are then managed by UniCredit's CAGP both to support internal operations and to publish them to other departments of UniCredit CFO and to the bank's legal division.

Working with DataSynapse guaranteed the minimum processing time for even the most complex of trade pricing and modeling calculations. GridServer has maximized computing efficiencies with UniCredit's CAGP now able to price more than two million deals in less than two hours. Before DataSynapse, this process took more than a day. These figures are then forwarded to the bank's accounting, finance and legal entity departments.

As one of the first Italian banks to leverage virtualization technology, UniCredit significantly reduced its operational risk and lessened redundancy in data generation. By having a unique financial model that is valued at a group level, UniCredit avoided any data inconsistencies provided by different groups within the bank.

Source: DataSynapse

Application Virtualization Improves Business Intelligence

Whether it is time critical customer analytics or the production of financials at quarter's end, today's enterprise is increasingly reliant on reporting and analysis of trusted and timely operational data to gain and maintain competitive advantage and comply with regulatory requirements.

Therefore, it is only natural that the enterprise is experiencing rapid growth in end user demand on mission critical BI and DI applications. Organizations simply can't keep up with the change in business requirements and user demand. BI and DI applications are already big and volatile consumers of compute resources, but the challenge is exacerbated by the rising volume of data that must be processed from a number of fragmented source systems.

To meet this growing demand, BI and DI applications are often massively over-provisioned for peak load. This approach is effective but inefficient – and expensive. Worse yet, this model may still constrain end users with performance and availability problems caused by under-provisioning. The practice of provisioning dedicated hardware is proving both costly and inefficient.

In this context, cost effective scaling of BI and DI applications for mission-critical performance and availability presents a growing and ongoing IT challenge. Overcoming these challenges through innovation will define an organizations competitive edge.

DataSynapse FabricServer™ decouples static and siloed BI and DI applications from dedicated data center infrastructure, aggregates and allocating server capacity in real-time based on business policy. With this policy-based automation, resources can be allocated to BI AND DI applications based on demand, schedule, and priority to improve quality of service (QoS).

In fact, BI and DI applications have an excellent general use case for sharing resources with one another, which is what FabricServer is all about.

The diagram on the next page shows the BI applications demanding and using more resources while people are running reports and performing analysis during business hours (9am to 5pm). At night (5pm to 9am), however, DI applications have the balance of resources to be used for data extraction, transformation, transportation and loading (ETL) in preparation for the next day of

reporting and analysis in the BI environment. This is an example of schedule-based policies that are used to address predictable fluctuation in demand.

But demand for IT resources is dynamic and usage patterns can be more volatile than they are predictable. Even so, service quality must be met cost effectively through automated service level management (SLM). In another example of policy-based automation, the diagram in Figure 1 depicts FabricServer monitoring the load on a BI AND DI application. When performance falls out of a pre-defined and normal operating range in the "After 5" load scenario, FabricServer will allocate more resources to BI AND DI. This is an example of a policy-based automation.

In this scenario, mission-critical BI AND DI applications get compute resources when they need them most. The benefits include improved Quality of Service (QoS), economics and agility. Document search and retrieval are faster, more accurate and based on freshest content. The organization can make better decisions faster, assuring their competitive advantage with greater confidence and reducing the cost of regulatory compliance.

5 Ways to Accelerate Performance of Business Intelligence Applications

Live and On Demand Webcast Series

Learn how you can accelerate Business Intelligence (BI) application performance and availability with application virtualization in this three part webcast series. Hear from the expert, Dan Kusnetzky, President and Principal Analyst from the Kusnetzky Group, as he explores real-customer issues surrounding business intelligence and the solution and benefits that application virtualization can bring to your enterprise.

Attendees will receive a complimentary analyst white paper associated with each webcast topic.

- Application Virtualization and Utopia featuring guest speakers from Cognos
- Unparalleled Performance: Harvesting the Power of Application Virtualization featuring guest speakers from Kusnetzky Group
- Application Failure is Not an Option: Achieving High Availability

To view a webcast, visit:

http://www.datasynapse.com/Webcast_BI_Series1_lp.aspx

DataSynapse has partnered with market-leading BI AND DI Independent Software Vendors (ISV) and delivered out-of-the-box FabricServer product integrations for core products in these ISV portfolios. FabricServer is a non-invasive technology, these software products require no code changes and function as normal. For example, FabricServer works in conjunction with an application's native load balancing functions.

There are two primary benefits to this approach. First, no retraining is required on the applications themselves when deploying them on FabricServer. Second, DataSynapse can greatly reduce time to market of new ISV product integrations in support of client-sponsored projects.

FabricServer for BI AND DI applications presents a high value, low complexity stand-alone project opportunity. The value is not limited to steady state, BI AND DI applications in

production. Because FabricServer automates the deployment of application code from a single point, rollout can be accelerated and time-to-value of major application upgrades and expansions can be reduced. Since production applications receive resources from a single resource pool, the need for extensive capacity planning can be eliminated.

With FabricServer, IT can adapt to changing application demand and cost-effectively align with the business in real-time through the supply of compute resources in the data center. Using FabricServer more broadly across an application portfolio improves the business value for virtualization geometrically. It is also a practical first step or compliment to other IT initiatives including server virtualization, SOA, business continuity and server/data center consolidation.

Source: DataSynapse

DataSynapse to Host Inaugural User Conference

Application Virtualization Pioneer to Host Event for Customers and Partners to Discuss Industry Trends and Technologies

DataSynapse will be hosting their inaugural user conference – ALIGN 2007 – on Sept. 17-19, 2007, in Miami. The conference will bring together DataSynapse executives and technical staff with partners, associates and customers, like Wachovia, UPS and LexisNexis, to discuss virtualization strategies and technologies to better align IT with business goals.

“Our clients are leveraging DataSynapse software to create a real-time infrastructure to better align IT with business demand,” said Peter Lee, CEO and co-founder of DataSynapse. “We’re pleased to provide this forum for peer interaction and product education. It’s an exciting milestone in our company history.”

Held at the Doral Gold & Spa Resort, ALIGN 2007 will be rich with business and technical content, including an array of exciting keynote speakers and sessions. The conference will include:

- Real-world case studies from leading global enterprises
- Presentations on the virtualization ecosystem and its impact on infrastructure operations
- A Partner Pavilion featuring virtualization experts, presentations and product demonstration from DataSynapse’s world-class technology partners
- Product demonstrations and training featuring the latest releases of DataSynapse’s flagship products, GridServer and FabricServer

For more information on ALIGN 2007, including registration dates, speaker information and session schedules, please visit: <http://www.datasynapse.com/user-conference-align-2007.aspx>.

Source: DataSynapse



About Us

DataSynapse is a global provider of application virtualization software. Our technology enables today's enterprise to implement a real-time infrastructure that better aligns IT with the business, delivering greater agility and improved service levels at radically lower costs.



Application virtualization decouples applications from their infrastructure, permitting dynamic provisioning and scaling of service execution. The company's flagship products, GridServer® and FabricServer™, enable enterprises to manage multiple applications in a shared environment and deliver capacity on demand. Both provide policy-based command and control to guarantee service levels for priority applications, but each targets different environments.

GridServer virtualizes application services, addressing compute- and data-intensive applications for which distribution and caching of the application workload enables significantly increased scale and accelerated performance.

FabricServer virtualizes application platforms, addressing application servers, 3rd-party ISV platforms, or internally developed application containers where dynamic provisioning and re-sizing of the platforms ensures high availability for priority services.

The company's worldwide production clients, partner ecosystem and best-of-breed solutions have been recognized by influential publications and industry groups, earning coveted accolades year after year, including: Inc. 500 fastest growing private companies, InfoWorld 100, NetworkWorld Enterprise All-Stars, Red Herring 100, AlwaysOn 100, JBoss Innovation Award, and others.

DataSynapse has regional HQ in New York City, London and Tokyo and can be found on the World Wide Web at www.datasynapse.com.

DataSynapse Fast Facts

- Peter Lee (CEO) and Jamie Bernardin (CTO) founded DataSynapse in 2000.
- With a customer base that spans 13 countries, DataSynapse has sales and support offices across the United States, Europe and Asia.
- DataSynapse counts 21 of the Fortune 50 Top Global Banks as clients.
- In 2006, DataSynapse was ranked 39 on the Inc. 500 lists of Fastest Growing Companies in America.
- In 2007, DataSynapse CTO, Jamie Bernardin was named as a part of the "CTO 25" and named one of the "10 Best Products for Next-Gen Network Infrastructures".
- Our client implementations consistently earn industry recognition for innovation, ROI, and excellence, with award wins from ComputerWorld, NetworkWorld, InfoWorld The Banker, FST and Financial-i.

Real-Time Infrastructure is published by DataSynapse. Editorial supplied by DataSynapse is independent of Gartner analysis. All Gartner research is © 2007 by Gartner, Inc. and/or its Affiliates. All rights reserved. All Gartner materials are used with Gartner's permission and in no way does the use or publication of Gartner research indicate Gartner's endorsement of DataSynapse's products and/or strategies. Reproduction and distribution of this publication in any form without prior written permission is forbidden. The information contained herein has been obtained from sources believed to be reliable. Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information. Gartner shall have no liability for errors, omissions or inadequacies in the information contained herein or for interpretations thereof. The reader assumes sole responsibility for the selection of these materials to achieve its intended results. The opinions expressed herein are subject to change without notice.