

An Introduction to DynaCenter™ Technology

IT Complexity Leads to Inflexibility and Underutilized Assets

IT environments are growing increasingly complex to administer as organizations add more servers, heterogeneous platforms, multiple operating systems and patch levels, multi-tiered storage solutions, and an ever-increasing number of mission-critical applications. This complexity increases administrative costs, requiring more administrators, more training, and more time to install, configure and maintain each server and its surrounding storage and network environments. Complexity also increases administrative errors, resulting in increased downtime and increased security risks.

Due to the cost, time, and risk associated with provisioning a server and its surrounding storage and network environments, production servers are rarely reallocated to meet short-term requirements. Once they have been deployed, they remain dedicated to an application. New requirements are met by adding new capacity. When provisioning an application, most organizations dedicate enough server capacity to handle the projected peak load for the application. But most of the time the application runs well below peak and its dedicated capacity remains lightly used, wasting rack space, electricity, and cooling, while placing an unnecessary burden on the data center's administration, maintenance, backup, and disaster recovery. But should load ever exceed projected levels, the application becomes overloaded, causing unacceptable performance or a system failure. This static, inflexible allocation of resources limits most data centers to utilization rates of just 5% – 15%¹ of their installed server capacity, causing organizations to waste significant time and money on their IT infrastructure and its management.

The DynaCenter Solution

Racemi's DynaCenter technology solves the problems of complexity and inflexibility through automation, converting IT resource allocation from a technical challenge into a real-time business decision. DynaCenter splits the provisioning and deployment of your resources into two levels of management. At the lower level, it captures your best practices for server, network, and storage administration into templates from which it can fully provision and deploy applications, freeing your technical staff from repetitive operations. At the higher level it enables self-service IT, allowing non-technical users, if properly authorized, to allocate and deploy resources based upon business-driven decisions. DynaCenter's policy engine can even allocate and deploy resources automatically using the configured rules and policies in response to server or application failures, or overload/underload conditions.

The data center efficiency problem has attracted a dizzying array of technologies and products all aimed at reducing administrative burden and increasing data center efficiency. DynaCenter is the only such "automated resource management" (ARM) product that:

- Works cross-platform on all resources – all blades, all servers, and all operating systems² as well as the surrounding network and storage infrastructure
- Can re-deploy hundreds of servers in minutes rather than hours or days
- Works on your existing assets with **no risk** to the currently deployed applications.
- Can be deployed in days rather than months.
- More than doubles the utilization rates of your data center's assets

¹ IDC and Gartner

² Intel, AMD, SPARC, PA-RISC and Power and PowerPC processors. Windows, Linux, Solaris X86, BSD, Solaris, AIX and HP-UX operating systems.

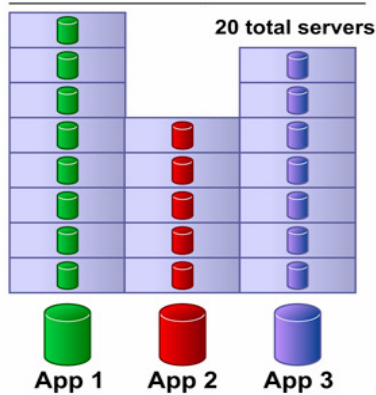
DynaCenter's Image Management

These advantages all come from DynaCenter's unique focus on **image management** rather than **server provisioning**. A server's software image consists of everything it would normally see within its file systems, including its kernel, operating system, applications, patches, data, etc. In order for a server to function, it must have a fully configured image. Traditionally this image has resided on the server's direct attached storage (DAS, i.e., local hard drive) and has been provisioned by administrators loading software from tape or CD-ROM, and then manually configuring the hardware, operating system and applications. Other ARM products automate this provisioning process, copying software over the network onto the server's DAS and then automatically configuring the server's hardware, operating system and applications. This is a huge improvement over the traditional, manual approach, and can dramatically reduce the time and administrative effort required to provision each server. But these other products only seek to automate the existing process. For them the acts of constructing the image, provisioning the server, and allocating the server for a specific application are all the same operation. Because they only maintain a single image per server, they must overwrite this image in order to repurpose the server for a different application. This can present problems when repurposing a server for short-term needs:

- Even with high-speed networks, it takes significant time to copy an entire image over the network onto the hard drive of a server. Repurposing a single server can take 20 minutes or more. 100 servers can take 20 hours or more.
- The original image will be destroyed when overwritten, losing any unique data or state.
- There is significant risk that you might not be able to return the server to its original purpose when needed. The management system or its infrastructure may be inoperable when needed. Even if available, the management system might not be able to exactly recreate the destroyed image.

Without DynaCenter

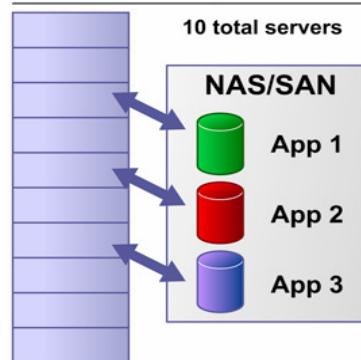
More servers required



Software images bound to local hard drives. Each server dedicated to a single application or service. Static server pools. Wasted resources.

With DynaCenter

Fewer servers required



Virtualized images stored in NAS or SAN. Servers free to run any ready image. Dynamic server allocation. Efficient use of resources.

NAS = Network Attached Storage
SAN = Storage Area Network
DAS = Direct Attached Storage
(i.e., local hard drive)

Image = a server's kernel, operating system, patches, applications, data, etc, which has been historically stored on DAS.

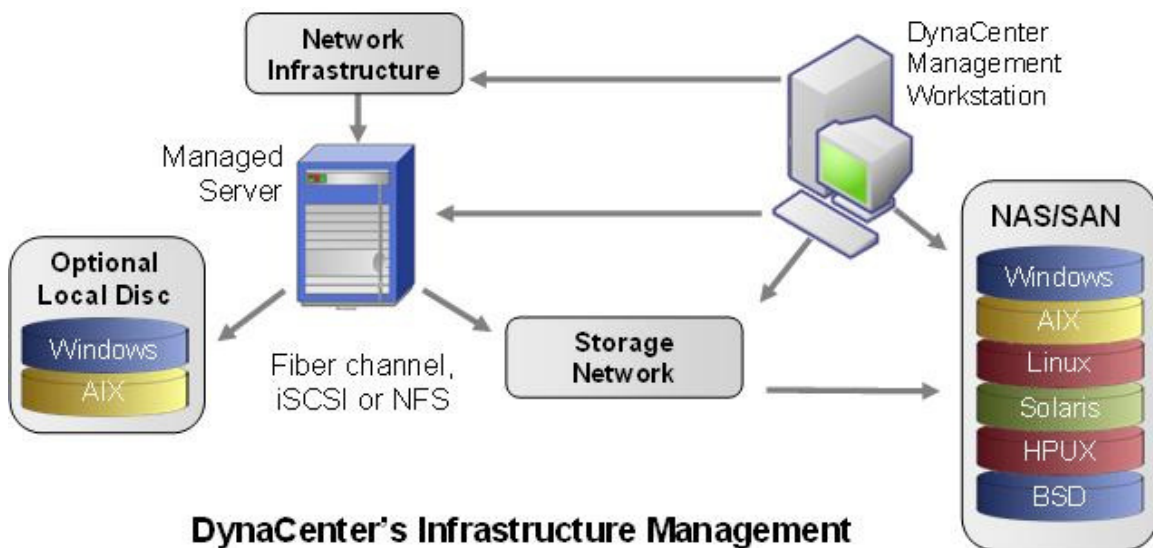
Unlike these other products, DynaCenter image management seeks not only to automate the process, but also to fundamentally improve it, as shown in the figure below

DynaCenter's Image Library

Like other ARM products, DynaCenter can directly provision a server and build out an image on the server's DAS. But it can also build out and manage images that reside in a centralized **image library** that is stored on network-attached storage (NAS) or storage area network (SAN). It can then direct a server to boot any appropriate image that is residing in the library³ or on the server's DAS. This separates the act of building an image for a server from the actual provisioning and deployment of the server. This has many critical benefits:

- An image can be built out in the library without disturbing the current operation of the server.
- Multiple images can be maintained for each server.
- Repurposing a server is almost instantaneous – as fast as a warm reboot – because there is no need to copy files onto the server's DAS. 100's of servers can be repurposed in minutes.
- Repurposing a server does not modify its previous image in any way, ensuring that the server can be quickly restored to its original purpose without risk⁴.
- An image in the library can be rapidly moved to any available server facilitating fail over.

As shown on the right side of the figure above, DynaCenter's use of the image library provides incredible flexibility in the deployment of your servers, freeing you from static server pools and allowing you to meet your processing requirements with far fewer resources. But to make this work in a production environment DynaCenter provides several additional features beyond its basic management of the image library. Some of these are shown in the figure below including:



- Automated management of the network and storage environments surrounding the server.
- Automated power control to turn servers on/off as needed.
- Remote access to a server's serial console.
- Security controls to ensure that a server can only access its authorized image.
- Security controls to restrict management features to properly authorized administrators.

³ Images in the library can be booted directly by a server using NFS, iSCSI, or fibre channel.

⁴ If the server's original image was on DAS, the server can be restored to that image with a simple reboot even during a failure of DynaCenter's management system.

- Accounting to track and report actual server use.
- A rich set of APIs to facilitate integration with other management systems within the data center.
- Full software life-cycle management for all images whether on DAS or in the image library. Install software and patches one time and have them automatically deployed across entire server pools or rolled back as needed.
- Image migration – move an image back/forth between the library and a server's DAS.
- Image virtualization – move an image from a server of one hardware configuration to another⁵.
- Shared, read-only image segments that can reduce total image size by 80% - 95%⁶.
- Complex image layouts combining both private and shared file systems on DAS, NAS, and SAN.
- Server and application monitoring to detect failures and current health and load conditions.
- Policy-based resource allocation triggered by a scheduled event, failure events, or real-time load and health conditions.

The Benefits Provided by DynaCenter

DynaCenter advanced image management provides the data center with a variety of important benefits including reduced administrative expenses, reduced risk, reclaimed resources, and new revenue streams.

Reduce Data Center Costs

- Asset utilization □, Capital costs □.
- Automation □, Operating costs □.
- Number of servers running at any time □, software licensing costs □.
- Convert fixed costs to variable with usage-based cost allocation.

Improve Service Response

- Automated resource allocation and provisioning – fully lights out.
- Improved application reliability. Provide automated fail over for all applications. Enforce service level agreements.
- “Instant” provisioning.

Drive New Revenue

- Rapid, automated application deployment.
- Enables variable costing. Provides the necessary accounting.

Reduce Security Risks

- Rules-based policy engine.
- Strict security controls.
- Read-only image segments provide resistance to hackers.

Bring Order From Chaos

- Automated, centralized resource management.
- Common, cross-platform interface.

⁵ For Unix and Linux images only, DynaCenter will reconfigure the image on the fly to account for changes in hardware. The target server must supply the minimal hardware required by the image, such as proper CPU-type.

⁶ Shared storage segments can be provided transparently to the operating system and applications in a Unix/Linux environment without any special installation or configuration. In the Windows environment shared storage segments require careful planning to isolate shared files on a separate “drive”.

Simpler, Faster, More Reliable Backups

- Central image library eliminates the need for network backups.
- Shared image segments reduce total backup volume 80% or more.

The Many Uses for DynaCenter

At first glance it is easy to understand how the flexibility and automation provided by DynaCenter might help you reduce your IT administrative costs and improve your resource utilization – if you had the luxury of enough time and resources to learn the technology and integrate it across your data center. But what about well-defined projects of limited scope that you can implement quickly and easily for an immediate return? Several examples are provided below. See Racemi's case studies for more detailed information.

Recover Your DR Assets

There are over \$50B of IT assets sitting idle in Disaster Recovery (DR) facilities⁷ waiting for something to fail. Some of them may be yours. Having made this investment, wouldn't you like to use these assets for other, non-critical applications while you wait for a failure, **without any risk** to your DR capabilities? With DynaCenter you can run alternate applications on your DR servers using DynaCenter's image library, leaving the servers and their primary applications untouched. When a failure does occur, you can always recover your original DR configuration within minutes, even if DynaCenter itself, the network, or the image library have failed. Using DynaCenter to reclaim your DR assets really is risk free.



Safely Use Your DR Assets for Alternate Applications When Not Needed for Fail Over

Lower Your DR Costs While Improving Fail Over

If you have more than one application requiring DR protection, you can use DynaCenter and a single, shared set of resources to provide fail over for all applications. With DynaCenter you can snapshot the exact configuration for each application and save it in the image library. When any application fails, DynaCenter can restore it onto the shared resources and have it operational within minutes. And while your DR resources are idle, waiting for a failure, you can use them for other, non-critical applications.

⁷ Tower Group and IDC.



Use One Set of DR Resources to Back Up Many Applications

In addition to reducing your DR hardware costs, DynaCenter can also reduce the associated software licensing costs. Because you are only running one application at a time on your DR servers, you only need one set of licenses that can be shared among the applications.

Image Libraries for Development, QA, Support, and Training Environments

Many development, QA, support, and training environments maintain large inventories of system images for testing and training purposes – different versions of a product or application environment, different test cases or scenarios, etc. This is often done by dedicating a server to each unique image or configuration – an expensive approach in terms of both hardware and administration. It is also risky, as any accident or error during normal use may corrupt an image, rendering it useless or perhaps corrupting test results.

Using DynaCenter's software management features you can snapshot individual software building blocks or configuration profiles and then assemble them on the fly to create any desired configuration. Using DynaCenter you can capture images directly from existing test systems and save them for future use within DynaCenter's image library. When they are needed again, they can be rapidly cloned and allocated to any available server, reducing the required number of test systems and facilitating automated test sequences across multiple images. As the master image is never used directly by a server, there is no chance of corruption.

Time-Share Applications

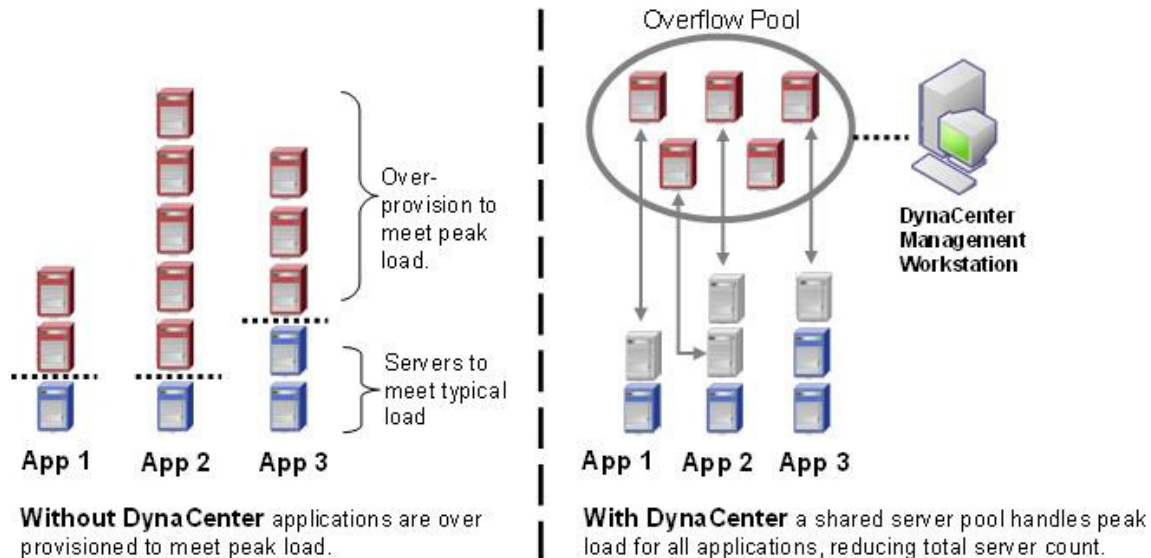
DynaCenter can rapidly and safely re-provision hundreds of servers in minutes, enabling you to reuse those resources that sit idle during specific times of the day, days of the week, etc. You can instantly change operating systems, applications, patch levels, etc. All you need is an alternate application that can run on those servers during the idle period. For example:

- Your desktop PCs running Windows sit idle in the evening and on the weekends. You can now switch them to a grid application running on Linux, and then back to Windows in the morning – or use them to create an automated QA test environment, etc.
- Your transaction processing systems sit idle when the market is closed. Switch them to a batch job or data mining application with full confidence that they will be processing transactions again when the market opens.

When time-sharing a set of applications, a server will only run one application at a time. This allows you to share one set of software licenses among the applications being time shared, further reducing your deployment costs for these alternate applications.

Server Pool Consolidation

If you have one or more application that experiences a varying load, DynaCenter can help you consolidate your server resources.



Using a Shared Overflow Pool to Reclaim Wasted Server Capacity

Rather than dedicate a full server pool, sized to meet expected peak load, to each application, you can dedicate just enough capacity to meet each application's typical load. DynaCenter can then manage a pool of shared servers, allocating each application the extra capacity it needs, only when it needs it. And when servers are idle within the shared pool, they can always be used for time-shared applications, as described above. Not only will this approach reduce the total number of servers that you must dedicate to your applications, it will eliminate application overloads caused by incorrect estimates of peak load. The more applications you are running and the more their load varies, the more you can save.

Understanding the Differences Between DynaCenter and Other Technologies

Provisioning Tools

Like DynaCenter, provisioning tools automate software installation and patch management for the servers in a data center, thereby reducing the administrative burden, improving reproducibility, and reducing errors. Most server vendors provide a provisioning system for their own platform, such as Sun's Jumpstart, Linux's Kickstart, and HP's Ignite. Cross-platform provisioning tools are available from a variety of vendors including Altiris, Marimba, Jareva (now Veritas), Novadigm (now HP), and Platespin.

All of these provisioning products work by managing an image on each server's DAS. Repurposing a server means overwriting the server's current image on DAS, which is both time-consuming and risky. It can take 20 minutes or more to repurpose a single server, hours or days to repurpose an entire server pool. And if the operation should fail, your servers can be left in an

inoperable state. Unlike DynaCenter, none of these products can manage images in a central image library on SAN/NAS and then boot a server directly from the library. They cannot install directly on your existing servers without disrupting current applications. They cannot instantly switch a server from one image to another. They cannot monitor applications and automatically repurpose servers based upon current conditions in the data center and the configured rules and policies. Nor can they provide fault detection with automated fail over.

Grid Technology

Like DynaCenter, grid technology helps you increase the utilization of your computers by allowing them to run alternate applications when they are not otherwise busy. Sophisticated grid systems, like those from Platform and DataSynapse, install on your existing compute resources without disturbing their current applications. When one of these "nodes" is not busy performing its primary application, the grid software checks in with a grid controller to see if it can help process any calculations. When the activity on its primary application picks up again, it disengages from the grid system. This works quite well for compute intensive applications. However, there are some important things to note about most grid products:

- They only work with a (currently) small number of compute-intensive applications that have been specifically modified to work in that vendor's specific grid environment.
- Most applications will never be suitable for the grid environment.
- All nodes in the grid must be connected via a high-speed network. This makes a grid unsuitable when the available compute resources are separated due to physical, administrative, or security concerns.
- Grid products only provision processor time and do not provision other resources around the node such as the network or storage infrastructures that are needed for the majority of applications.
- Because they do not provision network and storage access, grid products are unsuitable for applications that have security requirements in conflict with the server's primary application. For example, you wouldn't run a grid-based accounting application on your external web infrastructure. Any compromise of a web server might then compromise your accounting application.

DynaCenter and grid products are quite complementary. DynaCenter has been used successfully in several installations to dynamically switch server capacity between grid and non-grid applications. DynaCenter can scale out your grid capacity as servers become available, reconfiguring the network and storage environments surrounding the server to meet any security concerns. With DynaCenter managing a grid, it is safe to reuse your external web servers for your accounting grid application.

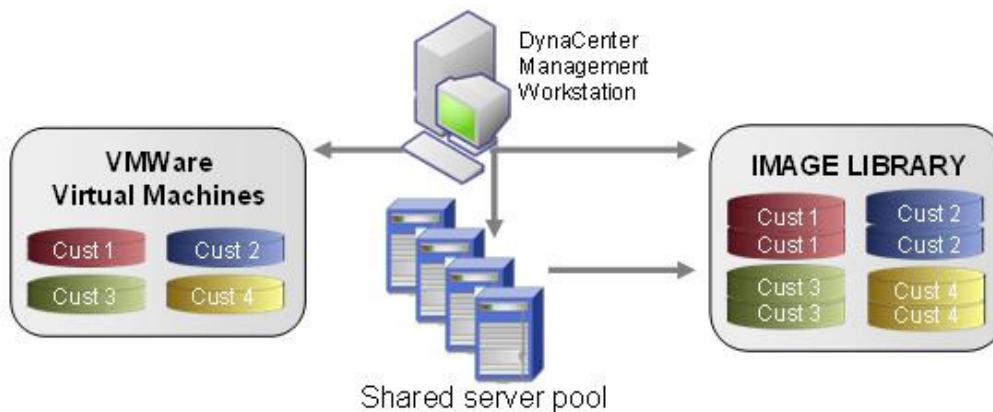
Server Virtualization Technology

Server virtualization technology allows a single server to run multiple server images at the same time, each securely isolated in a separate "virtual machine". This allows you to consolidate many lightly used servers onto a single piece of hardware, eliminating the extra servers and their associated costs. Server virtualization technology is provided on most high-end servers from vendors such as IBM, Sun, and HP. Vendors such as VMWare (now EMC) and Connectix (now Microsoft) provide server virtualization technology for industry-standard Intel-like servers running Windows or Linux.

While both DynaCenter and server virtualization products help you consolidate your servers, they are designed to consolidate different types of applications. With server virtualization products, all of the consolidated images are running at the same time. This means that each image requires a unique set of software licenses. It also means that the combined loads from the consolidated applications must not exceed the capacity of the shared server. With DynaCenter the consolidated images are not run concurrently. Thus they can share software licenses, and when running, can consume the entire capacity of a server. In other words:

- Server virtualization works best for consolidating lightly used applications with load requirements much smaller than the capacity of a server and that have similar security requirements such that they can share the same physical network and storage connections.
- DynaCenter works best for consolidating: a) applications that have variable load requirements that can exceed the capacity of a single server; b) applications that do not need to be online 24x7 and can be effectively time-shared; and c) applications that have incompatible security requirements requiring different configurations within the network and storage infrastructures.

Most data centers have a variety of consolidation needs, some ideal for DynaCenter, and some ideal for server virtualization. Often the two technologies can be used together to provide far better consolidation than either could provide on its own. Consider the case of an Application Service Provider (ASP) that provides a dedicated copy of its application for each customer. Each customer's copy needs to be available to the customer 24x7. But the load for each customer can vary widely, sometimes requiring multiple servers, sometimes going days without use. As shown in the figure below, a server virtualization product, such as VMWare, can be used to ensure that each customer's application is available to handle minimal load 24x7. But when appreciable load is encountered, DynaCenter can move the busy customer onto one or more dedicated servers.



Using DynaCenter with Server Virtualization to Obtain Optimal Consolidation

Single System Image

Single system image technology from vendors such as Virtuallron and XenSource creates a large, single working server from many smaller, individual computers ("nodes"). The resources of each node become available to all other nodes within the "server". They share processors and memory, and view a single, shared file system. Server virtualization technology is then used to segment this pool of resources into virtual machines, where a single virtual machine may now span multiple nodes depending upon its processing requirements.

This is an exciting technology that offers a great deal of promise for the future. However, at this time it:

- Requires an expensive Infiniband infrastructure to provide high-speed, low-latency communications between the nodes.
- Only works with Intel-like hardware.
- Requires a modified version of an operating system – most products currently support just Linux.

“Just-In-Time” or “Utility Computing” Computing Solutions

Products that provide true “just-in-time” or “utility computing” capabilities are those that compete most directly with Racemi’s DynaCenter. Like DynaCenter, these products add intelligence in the form of monitoring, analysis, and prediction on top of server, network, and storage provisioning so they can automatically allocate resources to applications when and where they are needed. Other vendors appear to be taking either of two approaches to just-in-time computing.

1. Companies such as RLX and IBM (through its acquisition of ThinkDynamics) use a software approach that could be called “smart” provisioning. They monitor the performance of servers and applications within the data center and apply trend analysis to predict future resource requirements for each application. While these products do improve resource utilization within the data center, they use standard provisioning technology to repurpose a server. Thus they are limited by the time and risk that this entails. Unlike DynaCenter, such products cannot react to unpredictable or rapidly developing conditions within the data center. They cannot rapidly repurpose servers to meet short-term needs. And they cannot install directly on your existing servers without disrupting your current applications.
2. Companies such as Egenera, Sun, and HP are taking a hardware-based approach. Egenera’s offering only works on its proprietary server blades. The offerings from Sun and HP only work on top of a fibre-channel/SAN based architecture. While these hardware-based offerings provide many of the features of DynaCenter, they cannot be deployed on your existing infrastructure and their initial build-out costs are quite high.

To Obtain Additional Information

As you have seen, DynaCenter’s combination of automated, intelligent provisioning and patent-pending Dynamic Server Allocation can greatly increase the utilization of your IT resources while reducing the cost of managing your data center. Please contact Racemi’s sales department for more detailed technical information on DynaCenter, detailed case studies, product data sheets, the current list of supported platforms, and pricing and availability information.

Racemi
380 Interstate North Parkway S.E.
Suite 250
Atlanta, GA 30339
800-457-1125
info@racemi.com
<http://www.racemi.com/>