

# Connection

Novell Connection Magazine // TOP TEN\_2010

## ARTICLES

PUBLISHED FOR  
*NOVELL CONNECTION*  
MAGAZINE.

**NOTE:** All content © 2010 by Novell, Inc. All rights reserved. For more information about Novell Connection Magazine or to obtain approval for bulk reprints, contact [editor@novell.com](mailto:editor@novell.com)

---

### Features

// **Hidden Gems - Part One**

(First featured in May 2010)

// **Mobilize Your Workforce with BlackBerry**

(First featured in October 2010)

// **Avoid Disaster in Disaster Recovery**

(First featured in August 2010)

// **Protecting Your Data with Novell Compliance Management Platform**

(First featured in September 2010)

// **Two Paths to Server Performance**

(First featured in July 2010)

// **Always in Sync**

(First featured in August 2010)

// **Population Automation**

(First featured in August 2010)

---

### Departments

// **From Microsoft Windows XP to Windows 7**

(First featured in July 2010)

// **Raise Your Workload IQ**

(First featured in October 2010)

// **Does Apple's iPad Really Mean Business?**

(First featured in July 2010)

# Hidden Gems - Part One

## Best Practices for Asset Management with Novell ZENworks

by Ken Baker

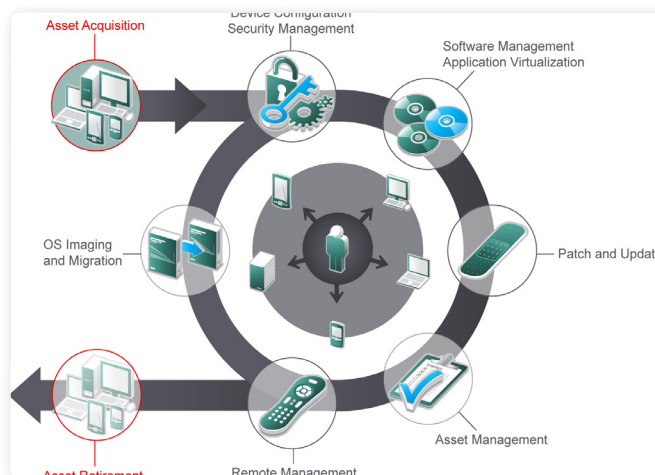
**The key to successful asset management lies largely in the processes you undertake. At Novell BrainShare 2010, the session entitled “Best Practices for IT Asset Management” discussed how Novell ZENworks Configuration Management and Novell ZENworks Asset Management can support, streamline and automate asset management best practices to reduce costs and simplify your management efforts.**

### > Pre-Deployment Management

If you're at all familiar with the ZENworks product line, you know that both [ZENworks Configuration Management](#) and [ZENworks Asset Management](#) can help you manage hardware throughout its different lifecycle stages. (See [Figure 1.](#)) However, what you might not know is how the products can help you in the acquisition stage—before you've even had a chance to deploy the ZENworks agent on your devices.

Presenter Christina Chamberlain, a Novell Technology Specialist, referred to this capability as a hidden gem that lets you add a list of devices into the ZENworks database by importing .CSV files. To take advantage of this import function, you need to know the serial number and hostnames of your devices. These values are used to reconcile the device to the record when the ZENworks agent is deployed. Depending on your procurement system, the serial number information might be available to you at time of purchase. If not, it might be available at time of receipt as part of a bill of lading or through the bar codes on the device packaging, which can allow the information to be scanned into a flat file before it's ever removed from the pallet on the receiving room floor.

Within the .CSV file device list, you'll also need a field column that contains the value WS\_1.0. This value is case sensitive and is required for the import to work properly. Also, the field order for your file needs to be the value WS\_1.0, the hostname of the device being imported, and then the device's serial number.



### Asset management lifecycle

**Figure 1:** Novell ZENworks Configuration Management and ZENworks Asset Management can help you manage hardware devices throughout their different lifecycle stages.

Once you have your .CSV file ready, you'll want to complete the following three steps in order to intelligently manage the import process:

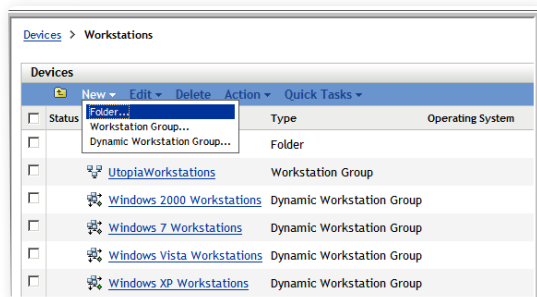
1. Create an Import Folder
2. Create a Registration Key
3. Add the Devices

### > Create an Import Folder

While you can import your new devices into any folder in the ZENworks database, management is simplified if you create a folder specifically for the imported devices. (See Figure 2.) This allows you to quickly see all your imported devices in one location. Also, if you create a rule to move the device to its appropriate site, group or department folder as soon as it has the ZENworks agent deployed on it, you only have to look inside your import folder to see which of your imported devices don't yet have an agent installed.

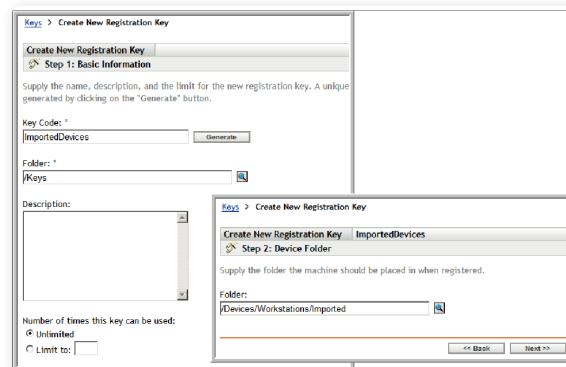
### > Create a Registration Key

The registration key is essentially a way of specifying what you want to do with each device record that is created when you import your device file list. For example, you can add it to a specific group, department or cost center. The registration key will also be used to tell the system what folder to move the record to, such as the import folder you created in step 1. (See Figure 3.)



**New folder dialog**

**Figure 2:** Management of your imported devices is simplified if you create a folder specifically for those devices.



**Registration key creation**

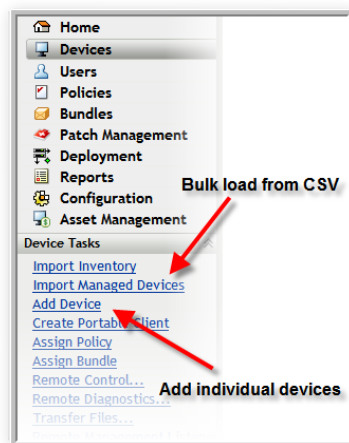
**Figure 3:** Creation of the registration key lets you specify what you want to do with each device record that is created when you bulk load import a device file list.

### > Add the Devices

On the left side of the ZENworks Control Center interface there are a number of powerful, yet often overlooked, task options. (See Figure 4.) One of these is the Import Managed Devices option, which enables you to import multiple devices using your device list file. (Note: There is also an Add

Device option that allows you to manually enter individual devices one at a time.) Once you have both your registration key and device list file correctly created, use the Import Managed Devices option to import your devices.

After you click this option, the interface will let you select your registration key and browse to your .CSV device list file. When importing is complete, you'll see all the new devices listed in the import folder you created. You can drill down into individual devices to see their unique hostnames and serial numbers. As mentioned before, this unique detail will enable you to automatically reconcile the imported device record with the actual device once it has the ZENworks agents installed.



#### **Bulk load from CSV**

**Figure 4:** The ability to bulk load import non-deployed devices into the ZENworks database is just one of a number of powerful task options available in ZENworks Configuration Management.

### > Tracking Device Life Changes

Once your devices are in the ZENworks database and deployed, you need to be able to keep track of their movement through your organization. Has a workstation moved from one location to another? Has it moved from one user to another? Has its internal hardware changed over time? Have changes been made to its software configuration? [ZENworks Configuration Management](#) and [ZENworks Asset Management](#) can help you track all of these types of changes.

Keeping track of a device's current user can be one of the biggest challenges you'll face. While some organizations rely on user input to keep this information up to date, you'll want to avoid relying on users for any inventory data entry. Even though ZENworks can simplify this process with the deployment of Collection Data Forms, the potential for human error often results in

---

**Whether it's the end of life, the beginning or somewhere in between, ZENworks Configuration Management provides the tools you need to manage your hardware assets throughout their entire lifecycle.**

---

too much time wasted by IT personnel having to scrub and clean that data. To avoid excessive data scrubbing, some organizations hire IT temps to perform regular inventory updates, but this still consumes valuable IT time and budget. A better option is to rely on the LDAP integration in both ZENworks Configuration Management and ZENworks Asset Management.

Within the ZENworks Control Center, you can create an LDAP Import Task that automatically and regularly imports demographic information (including device user names) for all your devices from either a Novell eDirectory or Microsoft Active Directory user source. To set up the LDAP Import Task, you select the Asset Inventory tab from the Configuration menu and then click New to launch the New LDAP Import Task Wizard. The wizard will ask you to specify an LDAP source that you have already defined as a user source in your Management Zone. It will also ask you to select and map the fields you want imported into the ZENworks database, as well as create the schedule that determines when and how often the task will be executed.

---

## **The “Retired” status in ZENworks Configuration Management or ZENworks Asset Management can help you address your beyond end-of-life management requirements.**

---

Just as you can schedule LDAP Import Tasks, you can also schedule regular inventory scans to determine the current status of your devices, such as Deployed, In Service or Retired. Inventory scans will report on changes within the device hardware, including internal memory, total disk space, free disk space and BIOS versions. Inventory scans can also report certain device software changes as they relate to anti-virus/anti-malware definition files, new applications, deleted applications, undesirable software and any deviations from your defined software standards. While all of these functions are available in ZENworks Configuration Management, you can obtain a greater breadth and depth of software reporting by taking advantage of ZENworks Asset Management.

From these inventory scans, the ZENworks product line provides a broad array of reports that can play a critical role in planning and budgeting. Both ZENworks Configuration Management and ZENworks Asset Management provide system readiness reports that can tell you which devices meet the performance hardware requirements for OS migrations or applications upgrades. ZENworks Asset Management is also able to report on trends and application usage. ([Read part 2 for more on this subject.](#))

It’s important to note that both ZENworks Configuration Management and ZENwork Asset Management, as well as ZENworks Patch Management, leverage the ZENworks Control Center as a common management console. Additionally, ZENworks Asset Management is included in the Enterprise Edition of ZENworks Configuration Management.

> **End of Life and Beyond**

New regulatory requirements have made end-of-life tracking and reporting more important than ever. In certain industries, you are responsible for every bit of data on a machine, even after its useful lifespan within your organization. So even if you have sold, donated or thrown away your retired machines, you might need to be able to certify that all the data (or the hard drive where the data resided) has been destroyed.

---

## **Novell ZENworks Configuration Management and Novell ZENworks Asset Management can support, streamline and automate asset management best practices to reduce costs and simplify your management efforts.**

---

The “Retired” status in ZENworks Configuration Management or ZENworks Asset Management can help you address your beyond end-of-life management requirements. With this in mind, you need to be aware that if you delete a device in the ZENworks Control Center, it completely removes the device and all its associated data from the ZENworks database. However, if you simply retire a device in the ZENworks Control Center, all of its associated information will remain accessible from the database, giving you reporting capabilities that can support compliance efforts. Another benefit of retiring a device is that all software applications associated with the retired device are excluded from software compliance reporting.

Whether it’s the end of life, the beginning or somewhere in between, ZENworks Configuration Management provides the tools you need to manage your hardware assets throughout their entire lifecycle. And by adding ZENworks Asset Management into the mix, you can better manage your software assets as well. To learn more about how ZENworks Asset Management supports the software-based asset management best practices discussed at BrainShare 2010, read part two of this article ([Hidden Gems – Part 2: Best Practices for Asset Management with Novell ZENworks Asset Management](#)).

---

**Online Resources** // Novell Connection Magazine

### **Learn More about Novell ZENworks**

- [Hidden Gems - Part 2](http://www.novell.com/connectionmagazine/2010/05/hidden_gems_part_two_one.html) (www.novell.com/connectionmagazine/2010/05/hidden\_gems\_part\_two\_one.html )
- [Novell ZENworks Asset Management](http://www.novell.com/products/zenworks/assetmanagement) (www.novell.com/products/zenworks/assetmanagement)
- [Novell ZENworks Configuration Management](http://www.novell.com/products/zenworks/configurationmanagement/) (www.novell.com/products/zenworks/configurationmanagement/)
- [Reviewer’s Guide](http://www.novell.com/docrep/2010/02/ZAM%20Reviewers%20Guide_en.pdf) (www.novell.com/docrep/2010/02/ZAM%20Reviewers%20Guide\_en.pdf)

# Mobilize Your Workforce with BlackBerry

Experience the Evolution of Mobility with BlackBerry Enterprise Server 5.0.1 for Novell GroupWise

by Ken Baker

---

**Business can happen any time at anyplace; you need to be just as productive on the go as you would be in the office. Research in Motion (RIM) in collaboration with Novell further enhances mobility for GroupWise users with the recent release of BlackBerry Enterprise Server V5.0.1 (SP1) for [Novell GroupWise 8](#).**

**“Our main goal is to fully, securely and efficiently mobilize GroupWise users,” says Slav Koziarski, Product Manager for Research In Motion. “Both RIM and Novell are committed to continue working closely together and provide our customers with high-quality products which address their mobility needs. Our product management and development teams have worked very hard to make this release the best to date.”**

## > More features for Mobile Users

Fully supporting BlackBerry smartphones running wide range of BlackBerry Handheld Operating Systems including version 5 and 6, BlackBerry Enterprise Server (BES) 5.0.1 for Novell GroupWise delivers an array of new features to help your GroupWise 8 users and administrators as well as your BlackBerry users be more productive while mobile. One of these new features allows you to synchronize multiple GroupWise address books with your BlackBerry smartphone, a feature that has been widely requested from the user community. In the past, you only had the ability to synchronize a single address book. But now if you have personal address books in addition to your business address book, you can synchronize them too, giving you convenient access to all your contacts.

---

**Both RIM and Novell are committed to continue working closely together and provide our customers with high-quality products which address their mobility needs.**

---

This version of BlackBerry Enterprise Server for Novell GroupWise also provides users with remote access from their BlackBerry devices to files stored on their office servers. (See [Figure 1](#).) Specifically, the “Files” feature allows you to securely retrieve, view, edit\* (Requires additional editing software such as Docs To Go) and e-mail documents stored behind your firewall on your Windows network file shares. So, if you’re running to a meeting and realize you need to look at

a certain file stored on your network, you can now access it from your BlackBerry smartphone. Or, if you are a frequent road warrior and need access to certain files while away from the office, you can leave those files on your network share and access them from your BlackBerry device as needed\* (minimum device code version 5 needed).

---



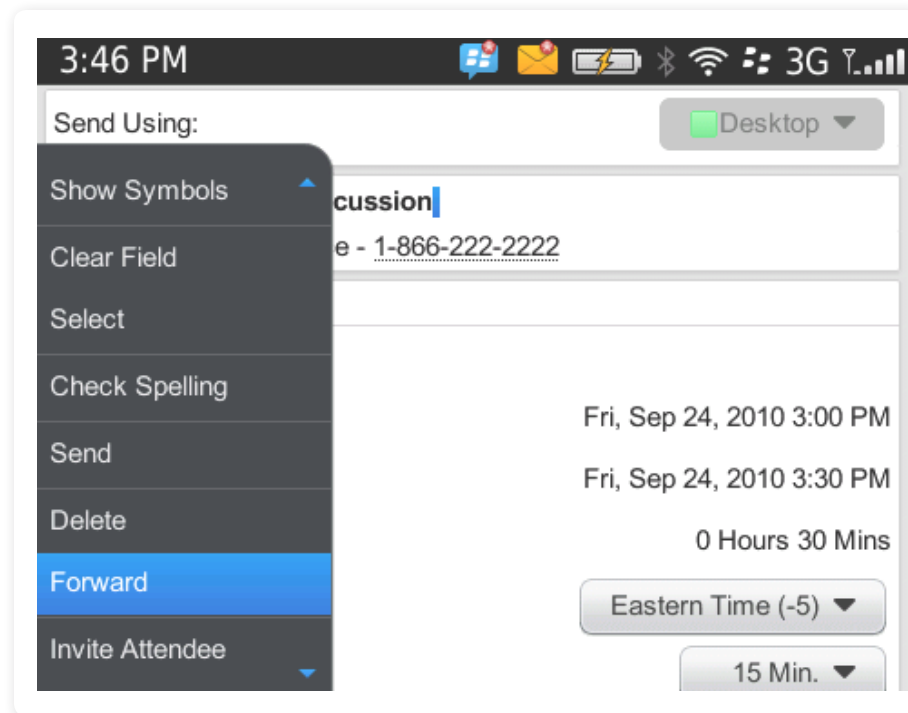
**Remote file access**

*Figure 1: BlackBerry Enterprise Server for Novell GroupWise lets you remotely access from your BlackBerry smartphones files stored behind your firewall on your Windows network file shares.*

The attachment file support in this latest release of BlackBerry Enterprise Server for GroupWise has been expanded as well. It now lets you open and view any word processing document, spreadsheet or presentation that you might receive as an Open Document Format (ODF) e-mail attachment on your BlackBerry. This includes documents created with OpenOffice.org applications.

Additionally, if you have the BlackBerry Presenter peripheral, you can run a whole PowerPoint or PDF presentation directly from your BlackBerry handheld. This reduces the need to carry your laptop around to meetings just so you can give a presentation. You can leave your laptop back at the office and just bring your BlackBerry.

BlackBerry Enterprise Server 5.0.1 for GroupWise now lets you forward calendar meeting invitations including those invitations that contain file attachments, directly from your BlackBerry smartphone. (See [Figure 2](#).) It also gives you the power to create, move, browse, rename and delete your GroupWise e-mail folders from your BlackBerry just as you can from your GroupWise client. All of these new features empower your mobile users, ensuring they have the right tools to drive business on the go. After all, there is no need to always have your laptop with you.



**Forwarding appointments**

*Figure 2: From your BlackBerry, you can forward calendar meeting invitations, including invitations that contain file attachments.*

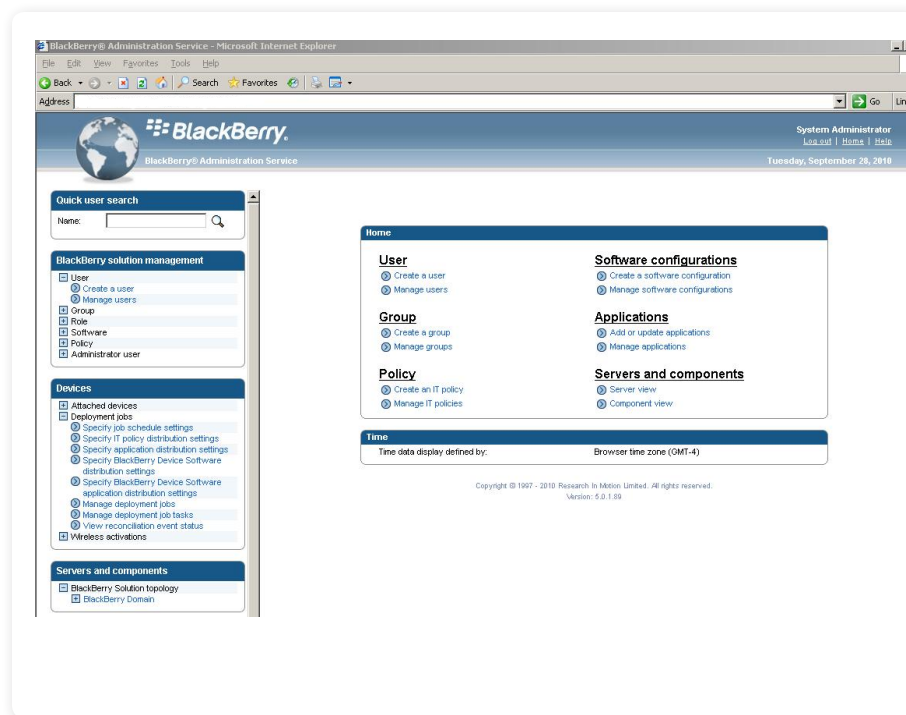
**> Simplifying Mobile Management**

The enhancements to BlackBerry Enterprise Server 5.0.1 for GroupWise not only address the wish lists of the BlackBerry and GroupWise user community, but this release also addresses several key requirements of the IT administrator community as well. As a primary example, the product is now much easier to install and manage. To help you get a feel for the ease of installation, you can take a virtual installation walkthrough by visiting this [tutorial](#) provided by RIM.

In terms of ease of administration, RIM has introduced the BlackBerry Administration Service (BAS), which offers dynamic, Web-based administration. BAS provides admins with centralized management of your BlackBerry Enterprise Server, letting you manage user accounts; assign user groups, administrative roles and software configurations; and apply IT policies to user accounts from any browser on your network rather than having to go to a specific server. (See [Figure 3](#).)

As a replacement for the BlackBerry Manager Console, the BlackBerry Administration Service has a new design and layout, but these changes are designed to make administration more intuitive and simple from a browser-based environment. To give you a feel for how simple it is to carry out certain administration tasks, you can visit [RIM's interactive tutorial](#) for the BlackBerry Administration Service.

RIM has also incorporated some new functions into the browser-based administration service. One of these is the addition of new IT policies that give you more control over what your users can and cannot do. For example, you can use these IT policy rules to manage security and behavior features of your BlackBerry devices in relation to encryption, user passwords and passphrases, protection of user data, and control of device resources (i.e., cameras and GPS capabilities) that are available to third-party devices. There are over 450 IT Policies in total.



### Administration interface

Figure 3: The BlackBerry Administration Service (BAS) simplifies management of your smartphones, user accounts and policies with centralized Web-based administration.

The BlackBerry Administration Service also lets you create roles for your administrator accounts, allowing you to assign different levels of access and permission to different administrators. It comes with pre-defined roles that determine what tasks administrators can perform. For example, it has pre-configured roles for security managers, enterprise-wide admins,

senior helpdesk engineers, junior helpdesk technicians, server admins, and even a role with limited permissions that allows individual users to manage their own devices. In addition to the pre-defined roles, you can create custom roles based on your unique needs. Roles can be assigned to individuals or groups.

To further simplify management of your BlackBerry devices, the new browser-based console also provides monitoring dashboards that give you instant visibility into performance and health metrics such as number of devices contacted recently, user statistics and the status of various BES components. In addition to at-a-glance statistics, you can drill down further into areas such as system health monitoring, which provides a unified, color-coded view of your BlackBerry environment as it relates to potential server or user issues. You can also monitor the health of your BlackBerry environment from a BlackBerry device by leveraging the new Web-based handheld dashboard.

---

## Now if you have personal address books in addition to your business address book, you can synchronize them too, giving you convenient access to all your contacts.

---

But the feature that will likely simplify administration of your BlackBerry environment the most is the new Blackberry Web Desktop Manager. This new tool lets you introduce self-service into Smartphone administration, allowing your users to address some of the most common device management tasks on their own. For example, it allows users to create activation passwords so they can activate their BlackBerry devices over the wireless network without having to contact your helpdesk. You can also give users the ability to use the Web Desktop Manager to configure how their address books and calendars are synchronized, as well backup and restore settings for data on their devices.

### > **Enhancing Mobile Management**

Another enhancement that will appeal to IT administrators is that BlackBerry Enterprise Server 5.0.1 for GroupWise now fully uses the SOAP protocol to interact with Novell GroupWise servers. This fundamental architecture change results in improved stability, increased performance and better scalability. It also eliminates the need to install the GroupWise client on your BES servers. In terms of improved scalability, testing indicates that a server running BlackBerry Enterprise Server for GroupWise can accommodate up to 500 users per server under a fairly heavy load. Capacity of previous BES versions varied based on certain environment configurations and therefore may have required additional BES servers to accommodate larger deployments. Upgrading to BES 5.0.1 will allow you to do some server consolidation—saving both hardware and licensing costs. And to help with that server consolidation, RIM provides the Enterprise Transporter tool, which allows you to effortlessly move users between BES servers without having to delete the user account on the original server and then reactivate it on the new server. Not only does this save you administrative effort and time, but the consolidation is completely transparent to the end user.

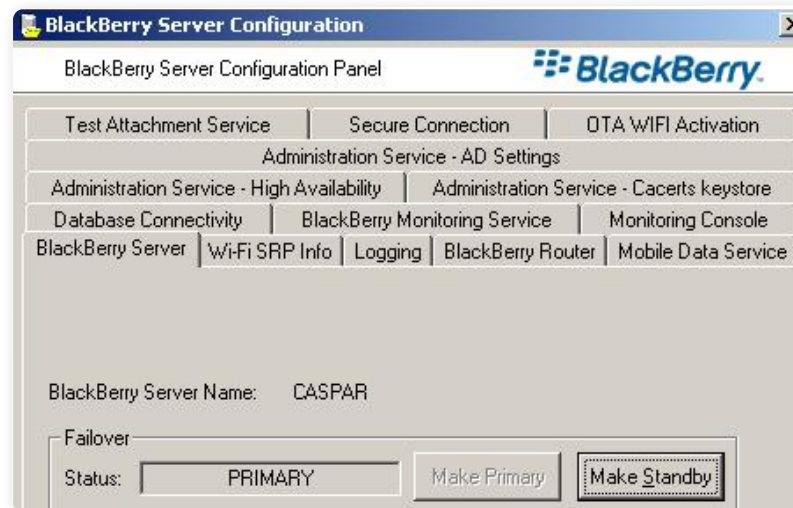
With the use of the SOAP protocol, you'll need to make sure that you have turned on SOAP at all the GroupWise Post Office Agents (POAs) where you have BlackBerry users. Additionally, if you use a proxy server when configuring the SOAP port, make sure your proxy server is transparent.

---

## The enhancements to BlackBerry Enterprise Server 5.0.1 for Novell GroupWise not only address the wish lists of the BlackBerry and GroupWise user community, but this release also addresses several key requirements of the IT administrator community as well.

---

BlackBerry Enterprise Server 5.0.1 for [Novell GroupWise](#) also has built-in high availability to help you reduce the risk of downtime. It lets you set up a standby BlackBerry Enterprise Server server that your primary BlackBerry Enterprise Server server can fail over to if certain customizable thresholds are reached. (See [Figure 4](#).) It also gives you the option to do load balancing between your primary and standby servers for some shared service components. Finally, you can perform a manual failover to your standby server to allow you to carry out planned maintenance tasks without affecting service to your BlackBerry smartphones.



### Creating standby server

*Figure 4:* You can create a standby that your primary server can fail over to if certain thresholds are reached.

**> Secure, Scalable and Reliable Mobile Access**

The commitment between Novell and RIM to further improve GroupWise mobility has led to the release of a BlackBerry Enterprise Server for GroupWise v 5.0 SP1 product that lets you give your mobile users secure, scalable and reliable access to their critical business information. It also simplifies and enhances your administrative efforts while delivering the system-wide scalability and availability you need. To learn more about how you can mobilize your GroupWise users with BlackBerry Enterprise Server for Novell GroupWise, visit <http://na.blackberry.com/eng/services/business/server/full/>.

---

**Online Resources** // Novell Connection Magazine

**Learn More about Novell GroupWise and Blackery Enterprise Server**

- [Novell GroupWise](http://www.novell.com/products/groupwise/) (http://www.novell.com/products/groupwise/)
  - [Walkthrough Tutorial](http://docs.blackberry.com/en/admin/deliverables/20949/index.html?name=Tutorial+-+BlackBerry+Enterprise+Server+for+Novell+GroupWise) (docs.blackberry.com/en/admin/deliverables/20949/index.html?name=Tutorial+-+BlackBerry+Enterprise+Server+for+Novell+GroupWise)
  - [Interact Tutorial](http://docs.blackberry.com/en/admin/deliverables/12167/index.html?name=Tutorial+-+BlackBerry+Administration+Service) (http://docs.blackberry.com/en/admin/deliverables/12167/index.html?name=Tutorial+-+BlackBerry+Administration+Service)
  - [Blackberry Enterprise Server and Novell GroupWise](http://us.blackberry.com/apps-software/business/server/full/) (http://us.blackberry.com/apps-software/business/server/full/)
-

# Avoid Disaster in Disaster Recovery

PlateSpin Forge and PlateSpin Protect Span the DR Gap for Both Windows and Linux, Including SUSE Linux Enterprise from Novell

by Jason Dea

---

**When you hear “disaster recovery,” if you picture expensive like-for-like infrastructures replicating your most essential servers, perhaps it’s time to rethink disaster recovery (DR). If you envision cumbersome backup servers running all night followed by tape cartridges traveling to a storage vault, again, it’s time to reconsider your DR strategy.**

If you are a Linux administrator and see yourself orchestrating your DR solution with scores of commands, it’s time you find a more efficient way to manage DR in today’s environment of shrinking resources and growing IT needs.

In each of these scenarios a gap exists between the cost and the effectiveness of disaster recovery. You get either cumbersome and affordable or expensive and effective—the choices administrators have been facing for years.

Novell designed [PlateSpin Forge](#) and [PlateSpin Protect](#) from the ground up to span this gap, and Windows administrators have been using them for some time. Novell has recently released the latest updates to these products to now protect environments running Linux and mixed Linux-and-Windows environments.

---

**If you see yourself orchestrating your disaster recovery solution with scores of commands, it’s time you find a more efficient way to manage disaster recovery.**

---

## > Spanning the Gap

A high-availability, or DR-by-duplication, strategy focuses on protecting workloads using local clusters, duplicate host sites or other like-for-like physical infrastructures. These strategies provide very low often near zero recovery time objective (RTO) and recovery point objective (RPO). These are the time it takes to move from a failure back to full production and the amount of data you lose in the process of recovery, respectively. However, even though RTO and RPO can be very low, these solutions can be very costly and extremely complex to configure and manage.

Traditional backup, on the other hand, eliminates the need for duplicate systems by storing backup data on tape or imaging devices, but acquiring new hardware, configuring it to the state of the original and restoring the data from the backup archive can take a long time, reducing effectiveness and increasing administrator aggravation.

The PlateSpin Forge and PlateSpin Protect products from Novell span the gap between cost and effectiveness with an entirely different way to think about DR.

The concept is quite simple, and the solutions are really equally easy to administer. PlateSpin Forge and PlateSpin Protect create a virtual machine (VM) copy of protected physical or virtual workloads and update the virtual machines at intervals you specify from 60 minutes to weekly or longer. When a server fails, you receive a notification and switch the users to this backup VM, almost immediately. The VM resumes the application load using its most recent update. And once the failed server is operational, you can either restore the VM from scratch to the server or do a reverse replication of just the changes since the failure occurred. The reverse replication would be

---

## **The PlateSpin Forge and PlateSpin Protect products from Novell span the gap between cost and effectiveness with an entirely different way to think about disaster recovery.**

---

appropriate when the original server can be repaired. If the original server configuration cannot be obtained, PlateSpin Forge and PlateSpin Protect have the flexibility to restore back to any x86 server from any vendor.

### **> PlateSpin Forge versus PlateSpin Protect**

Before we go further, you should understand that PlateSpin Forge and PlateSpin Protect are like two versions of the same product. They provide the same capabilities, use the same administration console and are managed the same way.

[PlateSpin Forge](#) is an all-in-one appliance that includes the storage, replication software, remote management interface and hypervisor. You just plug the appliance into your network, configure it, and it begins to protect up to 25 workloads. Novell designed PlateSpin Forge for medium-size enterprises and branch and field offices of larger enterprises.

[PlateSpin Protect](#) offers more flexibility in larger, more complex networks. It includes just the software the replication software and management interface. You install PlateSpin Protect on your own virtual infrastructure connected to your own storage system. You can purchase as many workload licenses as you want, giving you maximum flexibility.

With both products, you manage the entire DR system, regardless of the number of licenses and DR appliances, through a single graphical interface. You can also integrate these products into your existing DR infrastructure with an SDK that lets you configure a custom solution.

### **> How You Manage Disaster Recovery**

Now that you have a basic idea of what PlateSpin Forge and PlateSpin Protect are and how they might fit into your DR plans, you're probably most interested in what you have to do as an administrator. Let's take a look at how you monitor workloads, configure workloads, execute a failover, perform a failback, and conduct testing. We'll also look briefly at some of the reporting capabilities.



**Main Dashboard**

**Figure 1:** Management of the PlateSpin Forge and PlateSpin Protect DR protection begins in the Web-based console with the Dashboard.

**> Monitoring Workloads**

The PlateSpin Forge and PlateSpin Protect Web-based console begins with the Dashboard. (See Figure 1.) Other than the name in the top left corner, the management consoles are nearly identical for both products.

The Dashboard gives you an overall view of the workloads that are being protected: how many workloads are protected, whether any have failed and how many workloads are not protected.

Simple icons in the Workloads Summary show you the status of each workload. A green light means the workload is running normally. An orange light means there is an error in the configuration; perhaps you forgot to configure some part of the protection. A red light means the workload has failed.

The License Summary shows how many licenses are being used.

---

## Simple icons in the Workloads Summary show you the status of each workload.

---

The Storage pie chart shows the status of the available storage. PlateSpin Forge includes storage in the appliance, but you can also attach external, SAN or NAS storage.

On the right side of the screen, you can see a log of the past, current and upcoming events.

**> Configuring Workloads**

To configure a workload, you just need to tell PlateSpin Forge or PlateSpin Protect where the target workload is on the network by providing a host name or IP address. (See Figure 2.) You also provide the administrator credentials so the product can interact with the workload as needed. Then you identify the schedule for incremental replication based on how often the data changes and how much network bandwidth you want to use for replication.



**Configuring protection for a workload**

**Figure 2:** Inputting the workload hostname or IP address and its credentials gives PlateSpin Forge or PlateSpin Protect the information it needs to find and protect the workload.

One nice feature Novell added to these products is Smart Replication, which slices large batch replications into small portions that can be processed quickly without overloading the network.

A number of other configuration settings include whether you replicate by file or data block, how often you want to replicate, how often you want the devices to ping the workloads to ensure they are still functioning, whether to encrypt transfer data, the volumes that are protected and so on.

These are all menu driven and easy to follow. You can also set these configurations individually for each workload, or you can configure sets of workloads using templates.


**> Managing DR by Workload**

Novell designed these DR products to take the guesswork out of managing them. For example, in the Workloads tab, where you will spend the bulk of your day-to-day management, you can see that the console activates only the commands that are applicable to each workload. (See [Figure 3](#).) If you check on a workload that has failed, the Prepare for Failover and Run Failover commands become active.

This window also provides the other information you need about each workload, such as whether it's online, its protection tier, its replication status and so on.

**> When a Workload Fails**

If a workload or a server fails, you have two options: conduct an immediate failover, wherein the VM will boot up and begin serving the applicable users from the last snapshot, or start with the Prepare for Failover command. If there is a chance the failover is a false alarm, say someone simply tripped over a network cable, you should start with Prepare for Failover. With this command, the DR appliance stages the failover by configuring the virtual machine to take over for the failed workload but leaves it in a paused state. You can then choose whether to complete the failover by selecting the Failover or Cancel Failover commands.



**Workloads tab**

**Figure 3:** Workload commands are available dynamically based on the workload you highlight in the workload protection list.

## Novell designed these DR products to take the guesswork out of managing them.

### > Restoring the Production Server

This is where PlateSpin Forge and PlateSpin Protect really shine. Traditionally, once you've repaired or replaced a failed server, reinstalling the OS, applications and data can be a nightmare. However, with these products, all these steps are a one-click process. They can build the new workload from scratch, or in the case of a repaired machine, you can choose to restore just the changes that have occurred since the workload failed. This includes restoring plug-and-play drivers that have been updated or that are different because you're using new equipment PlateSpin Forge or PlateSpin Protect automatically install any necessary up-to-date drivers for Windows and Linux.



**Reports tab**

**Figure 4:** PlateSpin Forge and PlateSpin Protect maintain a history of activities and events and provide reports online, as PDFs or as exportable data for spreadsheets and other applications.

### > **Testing Your DR Environment**

If your DR system failed, now that would be a disaster. Testing is an important part of every DR strategy, and in the past this often required building an expensive and cumbersome parallel network just for testing. These DR products from Novell provide a Test Failover command that takes a snapshot of the VM and boots the snapshot to a virtual environment that you have defined as the test network. You don't boot the backup VM itself, and you don't affect the production network, so you can run this test anytime that is convenient for you.

### > **Tracking DR Performance**

PlateSpin Forge and PlateSpin Protect keep a history of all DR activities and events and measures performance such as how long it takes to run a job. They provide this data in several reports that you can view online, save as PDFs or export to applications such as spreadsheets. (See [Figure 4.](#))

### > **An Easier, Affordable DR Solution**

When one of your workloads fails, you no longer need to go through a cumbersome several-step process from retrieving a flat-file archive, sourcing a new server or repairing the existing server, redeploying the operating system and other software, and rebuilding the data files just to get your services back up and running. And you don't need to employ an expensive duplicate system. With PlateSpin Forge or PlateSpin Protect, you have a backup of the workload ready to move into place almost immediately, and it restores back to your production servers just as smoothly, with very little disruption to the users on your network.

---

#### **Additional Resources** // Novell Connection Magazine

##### **Learn More About PlateSpin**

- [PlateSpin Protect](http://www.novell.com/products/protect) (www.novell.com/products/protect)
- [PlateSpin Forge](http://www.novell.com/products/forge) (www.novell.com/products/forge)
- At the bottom of the Workloads tab in the PlateSpin Forge and PlateSpin Protect console are the Workload Commands. These buttons become active depending on the status of the workload you check. [Click here to see the full list online.](#)

# Protecting Your Data with Novell Compliance Management Platform

Integration of IAM and SIEM Crucial to Regulation Compliance

by Eric Harper

---

**Most organizations today are governed by one or more of the regulations directing the protection of personal information. These regulations, such as HIPAA or PCI DSS for example, were written to control the collection, storage, maintenance, distribution and disposal of private data. Most of these guidelines include somewhat vague mandates to “protect” or “restrict access to” customer, patient or member data. Consequently, many vendors have come forward to help organizations comply. And most modern identity and access management (IAM) products do a fine job of validating identity, provisioning resources and enforcing access roles.**

However, IAM covers only part of the rules. Another important aspect of these regulations involves data access auditing. Auditors want you to track what happened, when it happened and who did it. Again, a large number of security information and event management (SIEM) vendors are able to satisfy the audit-log requirements of these various rules, laws and regulations. And again, most SIEM products do a good job aggregating security data from throughout the organization.

## > Without Novell: Two Silos, No Communication

The result is two distinct sets of data: one set controls who has access to the organization's resources (through the IAM access policies) and another set shows who is accessing the organization's data (via the SIEM system). Unfortunately, they're usually not very good at talking to each other, causing all sorts of problems. Here's one example:

In January 2010, Lincoln National Corp., a financial services company based in Radnor, PA

---

**Lincoln National lacked a system that would have noticed whenever two different people logged in with the same user name at the same time. If they had such a system, security personnel could have been notified, those users sharing credentials could have been identified and the policy violation could have been rectified. Because there was no integration between the IAM and SIEM systems, the policy violation went on for eight years.**

---

disclosed a security vulnerability that may have leaked the personal data of 1.2 million customers. An investigation revealed that some employees of Lincoln National and another one of its subsidiaries, Lincoln Financial Advisors, were using shared user names and passwords to access the portfolio information management system. Six shared user names and passwords, which were created as early as 2002, were found.

Obviously, sharing user names and passwords was a violation of Lincoln National's security policy. But they lacked a system that would have noticed whenever two different people logged in with the same username at the same time. If they had such a system, security personnel could have been notified, those users sharing credentials could have been identified, and the policy violation could have been rectified. Because there was no integration between the IAM and SIEM systems, the policy violation went on for eight years.

And here's the real scary part. The vulnerability was discovered in August of 2009 (five months before it was disclosed), but not by Lincoln National! Someone sent an anonymous tip to the Financial Industry Regulatory Authority (FINRA) who notified Lincoln National. A forensic security company was hired to investigate, and they're the ones who found the violation. Unfortunately, it's fairly common for an outside party to discover security problems like this.

According to the "2010 Data Breach Investigations Report" from Verizon Business, while 86 percent of data breach victims had evidence of the breach in their audit logs, 61 percent of victims didn't uncover the breach themselves—they were notified by a third party! As the report states, "Verizon's past research consistently finds that breaches are not found by the victim organization but by an outside party." How'd you like to be the one who got that call?

---

**Info** // Novell Connection Magazine  
**SO LOGICAL, YET SO RARE**

But don't all vendors' solutions integrate IAM and SIEM technology? Unfortunately, no. Most contemporary applications, which are called Compliance Management Systems, simply write potentially interesting events to a log file. When compliance to a regulation must be documented, someone must read, digest and extract the data in these files. Most vendors will tell you they have a tightly-integrated system, but most are IAM solutions with a SIEM system tacked on, leaving customers to build the integration themselves or pay consultants to do it for them. And those "integrations" are definitely not real-time. The Novell Compliance Management Platform is more than just a bundle of products. It's the marriage of technologies. It helps you close the gap between what's supposed to happen and what's actually happening.

---

Not only do organizations regularly fail to discover evidence of breaches in their own audit logs, but the length of time needed for a third party to discover the breach is inordinately long. The Verizon Data Breach report notes that fully 70 percent of breaches go undetected for months or more. In fact, "Over the last two years, the amount of time between the compromise of data and discovery of the breach has been one of the more talked about aspects of this report. It is not without reason; this is where the real damage is done in most breaches. That a breach occurred is bad enough but when attackers are allowed to capture and exfiltrate data for months without the victim's knowledge, bad gets much worse."

If only Lincoln National had a solution that integrated their IAM and SIEM systems in real time—a system that constantly correlated identity access and policy information, as the events happened, across the entire enterprise. With such a system in place, if anomalous activity occurs, the proper people could be immediately notified—not months after the damage has been done,

**The Novell Compliance Management Platform can tell you which users have been provisioned for a particular application, which employees are actually using the application, when they use it and what they do within the application. Only the Novell Compliance Management Platform can monitor those activities, not just for audit purposes, but to intervene—with remedies—at the time the activity is occurring.**

---

and not by a third party. If Lincoln National had a system like the Novell Compliance Management Platform, they could have avoided the embarrassing public disclosure and regulatory admonishments.

If Lincoln National had the [Novell Compliance Management Platform](#), it could have uncovered the sharing of user names and passwords when that activity first occurred—eight years before the bank became aware of it.

The Novell Compliance Management Platform can tell you which users have been provisioned for a particular application, which employees are actually using the application, when they use it and what they do within the application. Only the Novell Compliance Management Platform can monitor those activities, not just for audit purposes, but to intervene—with remedies—at the time the activity is occurring.

You may have also heard about the case of France's second largest bank, Société Générale. In 2008, they reported that "rogue" trader Jerome Kerviel had misappropriated over US\$7 billion—the single largest fraudulent act ever in the securities industry. Apparently, Kerviel built-up entitlements as he moved from one position to another, and from one department to another. Société Générale had policies in place prohibiting this accretion of entitlements. These policies specifically forbid someone with one authorization (such as invoice approval) from having other authorizations deemed to conflict (such as check signing). But Kerviel didn't simply acquire authorizations for his own account, he also tapped into accounts shared among traders (and others) in violation of the bank's policies.

### > Industry-Leading Technologies

Novell has been in the identity and security business for over a decade. In that time, they've built a host of technologies—such as Novell Identity Manager, Novell Access Manager and Novell Sentinel—that are considered industry-leading technologies. Novell is positioned in the Leader's Quadrant of Gartner Inc.'s Magic Quadrant for User Provisioning, Magic Quadrant for Web Access Management and, most recently, its Magic Quadrant for Security Information and Event Management.<sup>1</sup>

**> Conclusion**

Examples such as the data breach at Lincoln National and the fraud at Société Générale show how companies continue to struggle with issues of policy compliance. Novell delivers a platform that provides a real-time, enterprise-wide view of the enterprise to mitigate the risk posed by internal and external threats and, ultimately, to ensure an organization's image, brand and reputation are safe.

The Novell Compliance Management Platform combines powerful technology with documented best practices to provide the only real comprehensive approach to policy compliance. To learn more about the Novell Compliance Management Platform and how it can help organizations bolster security, go to: [http://www.novell.com/promo/home/integrated\\_identity.html?nov\\_gaevent=Homepage|Banner|Integrated\\_identity](http://www.novell.com/promo/home/integrated_identity.html?nov_gaevent=Homepage|Banner|Integrated_identity).

---

**Online Resources** // Novell Connection Magazine

**Learn More about Novell Compliance Management Platform**

- [Novell Compliance Management Platform](http://www.novell.com/compliancemanagement) (http://www.novell.com/compliancemanagement)

---

1The Magic Quadrants are copyrighted 2009 and 2010 by Gartner, Inc. and are reused with permission. The Magic Quadrant is a graphical representation of a marketplace at and for a specific time period. It depicts Gartner's analysis of how certain vendors measure against criteria for that marketplace, as defined by Gartner. Gartner does not endorse any vendor, product or service depicted in the Magic Quadrant, and does not advise technology users to select only those vendors placed in the "Leaders" quadrant. The Magic Quadrant is intended solely as a research tool, and is not meant to be a specific guide to action. Gartner disclaims all warranties, express or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.

# Two Paths to Server Performance

## I/O Scheduler and File System Selection Can Boost SUSE Linux Enterprise Server Performance

By Matthias G. Eckermann and Bill Tobey

---

**The Novell approach to assembling a SUSE Linux server distribution has always been to provide a wide range of the best packages and tools available from the community. Our goal is to give IT organizations the most flexible and versatile resource set for configuring and optimizing high-performance servers for a complete range of data center applications.**

**This article will explore two often-overlooked areas where [SUSE Linux Enterprise Server](#) provides multiple options that administrators can exploit to enhance server performance: the I/O scheduler and the file system.**

### > Meet Your I/O Scheduler

The I/O scheduler is the part of the kernel that handles read / write access to block storage devices—a USB stick, local disk, NAS filer, SAN, network file system and any other storage environment that holds data in blocks. A scheduler queues and sequences the execution of read-write requests in order to manage mechanical latency (the seek time related to head travel around the disk) and optimize data delivery performance. Its bag of tricks includes three techniques for manipulating the request queue:

- **Request merging** – Requests for data in adjacent blocks can be combined to improve throughput by reducing both seek time and the total number of syscalls required to service a request.
- **Directional (elevator) reordering** – Requests can be reordered based on location, to maintain head movement in one direction for as long as possible, using the same control methodology as an elevator to avoid service starvation at the disk peripheries.
- **Priority reordering** – Requests can be sequenced according to various priority schemes, such as a start-of-execution deadline assigned to each request at time of receipt.

### > The Four Types of Linux I/O Schedulers

There are four types of Linux I/O schedulers, each of which implements the basic sequencing techniques in different ways and combinations, providing significant variations in I/O performance with different application workloads.

**The NOOP scheduler** is the simplest of all Linux I/O schedulers. It merges requests to improve throughput but otherwise attempts no other performance optimization. All requests go into a single unprioritized first-in, first-out queue for execution. It's ideal for storage environments with extensive caching, and those with alternate scheduling mechanisms—a storage area network with multipath access through a switched interconnect, for instance, or virtual machines, where the hypervisor provides I/O backend. It's also a good choice for systems with solid-state storage, where there is no mechanical latency to be managed.

To activate the NOOP I/O scheduler for use with all applications and storage devices, edit your boot loader configuration settings to pass the kernel parameter: [elevator=noop](#).

**The Deadline scheduler** applies a service deadline to each incoming request. This sets a cap on per-request latency and ensures good disk throughput. Service queues are prioritized by deadline expiration, making this a good choice for real-time applications, databases and other disk-intensive applications. To activate the Deadline I/O scheduler for use with all applications and storage devices, edit your boot loader configuration settings to pass the kernel parameter: [elevator=deadline](#).

**The Anticipatory scheduler** does exactly as its name implies. It anticipates that a completed I/O request will be followed by additional requests for adjacent blocks. After completing a read or write, it waits a few milliseconds for subsequent nearby requests before moving on to the next queue item. Service queues are prioritized for proximity, following a strategy that can maximize disk throughput at the risk of a slight increase in latency.

The Anticipatory scheduler delivers best performance with Web and file servers, and desktops with single IDE/SATA disks. It is the default scheduler in the mainline Linux kernel, and can be activated by editing the boot loader configuration file to pass the kernel parameter: [elevator=as](#).

**The Completely Fair Queuing (CFQ) scheduler** provides a good compromise between throughput and latency by treating all competing processes even-handedly. Each process is given a separate request queue and a dedicated time slice of disk access. CFQ provides the minimal worst-case latency on most reads and writes, making it suitable for a wide range of applications, particularly multi-user systems. Because of our unique desktop-to-data center strategy, CFQ is the default I/O scheduler in SUSE Linux Enterprise Server 11. It can be activated by editing the boot loader configuration file to pass the kernel parameter: [elevator=cfq](#).

### > Making Per-Device I/O Scheduler Assignments

If you have multiple applications running on a server, using different storage environments, it's possible to make separate I/O scheduler assignments to optimize the performance of each application-storage pair. These assignments can even be changed in production. You can check the I/O scheduler setting for individual storage devices with the following shell command: [/sys/block/\\*DEV\\*/queue/iosched](#).

If desired, you can then re-set the I/O scheduler assignment for each device using this command: [echo SCHEDNAME > /sys/block/\\*DEV\\*/queue/scheduler](#).

### > Integrity, Performance and the Barrier In Between

Barriers are a feature the kernel's block I/O subsystem makes available to journaling file systems to protect data integrity. A barrier request temporarily locks the I/O scheduler's execution queue, ensuring that a sequence of journal write requests are securely committed to physical media before any subsequent requests are served. Barriers protect metadata and ensure file system integrity in the event of a system crash, but they do so at the expense of a noticeable performance penalty. Novell assumes a higher value for data integrity than performance in enterprise computing

environments, so barrier support is switched on by default in SUSE Linux Enterprise Server. It can be turned off to improve performance, but only by a knowledgeable administrator prepared to assume the risk.

- With reiserFS you can enable / disable barriers using the mount options: [barrier=flush](#) or [barrier=none](#).
- With ext3 you can enable / disable barriers using the mount options: [barriers=1](#) or [barriers=0](#).
- With XFS you can enable / disable barriers using the mount options: [barrier](#) or [nobarrier](#).

### > **File System Selection for Server Performance**

Another set-up decision that can significantly affect server performance is the choice of file systems. As is the case with I/O schedulers, [SUSE Linux Enterprise Server](#) ships with a number of file system alternatives, allowing administrators to match file systems and application workloads for optimum performance. Here are a few guidelines for making the right performance pick.

- **Choose ReiserFS for small files** – ReiserFS is best suited for applications that generate lots of small files, including mail, NFS and database servers, and for applications that use synchronous I/O.
- **Choose ext3 for small file systems** – Ext3, the default file system in SUSE Linux Enterprise Server 11, is best suited for small file systems of 100 gigabytes or less. If you're planning to use ext3 with large numbers of files in a single directory, you should consider enabling btree support. This can be accomplished with the shell command: `# mkfs.ext3 -O dir_index`. Note that btree support is enabled by default in version 11 SP1.
- **Choose XFS for large files and streaming media** – XFS is an excellent choice for large files and medium to very large file systems (up to 16 terabytes on 32-bit systems, or a theoretical 8 exabytes on 64-bit systems). Its low latency transfer characteristics also make it an ideal selection for streaming media applications. SUSE Linux Enterprise has supported XFS since version 8, and Novell is working closely with SGI to optimize its performance with future releases of SUSE Linux Enterprise Server. It merits consideration for any file system likely to exceed 100 gigabytes, unless other factors (e.g. many small files) dictate another choice.

XFS also offers a number of special features that can be particularly useful, including dump / restore for backup and recovery, online file system check, and online defragmentation.

- **Choose OCFS2 for cluster performance or high availability** – Oracle Cluster File System 2 (OCFS2) is a POSIX-compliant shared-disk cluster file system for Linux that is developed by the community under GPL. Because it provides local file system semantics, OCFS2 can be used with any application. Cluster-aware applications can leverage its parallel I/O support for higher performance, other applications can leverage its multi-node support to achieve higher availability through automated failover.

- **Consider btrfs for the future** – Btrfs is a new, copy on write file system for Linux aimed at bringing additional enterprise class file system features to the Linux kernel. Initially developed by Oracle, btrfs is licensed under the GPL and has quickly been adopted by the community. Long-awaited features include integrated volume management, copy on write, writable snapshots (and snapshots of snapshots), extents, dynamic inode allocation, checksums on data and metadata, online file system check and defragmentation, and integrated multiple device support.

Btrfs is still under intense development, but is included as a technology preview in SUSE Linux Enterprise Server 11 SP1.

### > **Measuring I/O Scheduler and File System Performance**

Once you've made your I/O scheduler and file system selections, there are many tools available to measure your configured system's performance. A few favorites include:

- **Bonnie** is a simple but useful tool that provides a variety of benchmarks on the speed of your file system, OS caching, the underlying device and your libc. It is supported in the SUSE Linux Enterprise Server distribution.
- **fiio** is an I/O tool meant to be used both for benchmark and stress/hardware verification. It has support for 13 different types of I/O engines. It is available at: <http://freshmeat.net/projects/fiio/>. Packages for SUSE Linux Enterprise Server 11 are available at: [http://download.opensuse.org/repositories/benchmark/SLE\\_11/](http://download.opensuse.org/repositories/benchmark/SLE_11/).
- **IOzone** is a file system benchmark tool that generates and measures a variety of file operations. iozone has been ported to many machines and runs under many operating systems. It is available at: <http://www.iozone.org/> and [http://download.opensuse.org/repositories/home:/mge1512:/benchmarking/SLE\\_11/](http://download.opensuse.org/repositories/home:/mge1512:/benchmarking/SLE_11/).

### > **Experiment to Find Your Optimum Configuration**

I/O scheduler and file system selection can have major effects on the performance of SUSE Linux Enterprise Servers. We strongly recommend that you experiment with different configurations to gain first hand experience. And watch for our upcoming article on cgroups for kernel resource management.

---

#### **Online Resources** // Novell Connection Magazine

#### **Learn More about SUSE Linux Enterprise Server**

- [SUSE Linux Enterprise Server](http://www.novell.com/products/server/) (http://www.novell.com/products/server/)
- [SUSE Linux Enterprise Server 11 SP1 Download](http://www.novell.com/promo/suse/sle11sp1.html) (http://www.novell.com/promo/suse/sle11sp1.html)

# Always in Sync

## A Deep Dive into Novell Data Synchronizer

by Ken Baker

---

**As you read in the June issue of Novell Connection Magazine (See [Higher Levels of Collaboration](#)), Novell Data Synchronizer keeps applications and mobile devices constantly in sync so users can always have access to the data they need. That article talked about the product's mobility pack, its many-to-many synchronize engine and how to configure connectors. This article takes a deeper technical dive into the way connectors facilitate synchronization, as well as a few other supporting architectural elements.**

### > Synchronization Channels and Filters

The majority of the intelligence in [Novell Data Synchronizer](#) resides in its connectors. Written to application-specific or standard APIs, the connectors enable Data Synchronizer to understand how and what needs to be done with data to be synchronized. For example, the SharePoint connector uses the SharePoint Web Services API to access data and push it through the system. The Salesforce.com connector uses the Salesforce.com Web services API. Similarly, the GroupWise connector utilizes the SOAP interface to synchronize e-mail, tasks, calendar data and contact information. The use of standard or native APIs allows each connector to be built specifically for its target application.

Each connector in Novell Data Synchronizer is comprised of two main components: channels and filters. Channels conduct the flow of data from the synchronization engine through the connectors to the connected systems. Filters act upon the data—translating and manipulating the data in a prescribed manner.

At a very high level, channels and filters work together in the following fashion:

1. As data moves through the Novell Data Synchronizer channels between two or more applications, the Novell Data Synchronizer connectors act as code converters and data filters.
2. The data flows through the Novell Data Synchronizer system in the form of XML files that are either in an application-specific format or an application-neutral format.
3. As data moves from the application toward the synchronization engine, it will initially be in an application-specific XML file format.
4. As the application-specific XML file travels through the channel, the connector filter will translate the XML file into an application-neutral format.
5. After it's translated into an application-neutral format, the XML file is stored in the synchronization engine database.
6. The application-neutral XML files stored in the synchronization engine database can be consumed by different application-specific connectors (i.e., GroupWise, SugarCRM, Sharepoint, Salesforce.com), which in turn can translate them into the appropriate application-specific format as they transfer them to their associated application.

If you're using Novell identity and security products, you'll likely recognize that this design is very similar to that of Novell Identity Manager. While Novell Identity Manager synchronizes identities from different sources and targets, Novell Data Synchronizer synchronizes data in a similar fashion. These similarities are perhaps most notable in the use of a two-channel system with a source channel for outgoing events and a SINK channel for incoming events. In Novell Data Synchronizer,

these channels exist in both the synchronization engine and the individual connectors.

The source channel in the synchronization engine is responsible for taking item or event data from its cache and passing it to the connector. After the connector translates the data into the appropriate application-specific format, the connector's SINK channel will pass that information to the application. For data coming from an application to the engine, the connector's source channel has responsibility for taking data from the application, while the engine's SINK channel will place the application-neutral formatted data into its cache.

> **How to Synchronize**

---

## Written to application-specific or standard APIs, the connectors enable Data Synchronizer to understand how and what needs to be done with data to be synchronized.

---

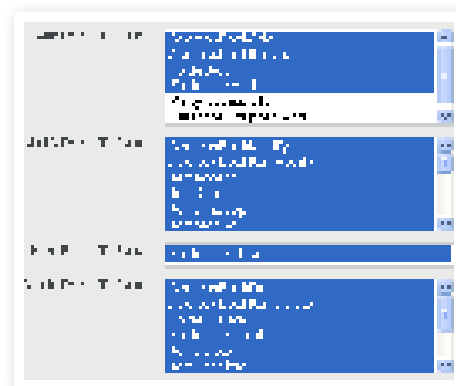
In addition to having translation filters, connectors have synchronization filters. Depending on application needs and purpose, each application will have different types of data that it will allow to be synchronized. Some of the common data types (referred to as data items) include e-mail messages, appointments, tasks, notes, address books, contacts and folders. Connectors will not necessarily support all the same types of data items. In other words, data item synchronization will vary based on the functionality and needs of the application.

In addition to determining what data types are synchronized, filters have responsibility for managing the synchronization of changes that occur to data items within a connected application, system or device. These changes (known as events) might be represented by actions such as an item addition, modification, move or deletion.

As data items and events pass through a connector, various filters applied to the data determine what and how the data is synchronized. Depending on the connector design, you can often use filters to determine which data items and events you want to synchronize.

Not only can you use filters to screen for certain content or fields, but you can use them to transform one data type into another data type, such as turning task data into e-mail data. Filter behavior can also be modified based on certain triggers. For example, you can have a trigger that causes a filter to treat a data item differently based on where that data was originally stored on the source application.

You configure a connector's filters through the Novell Data Synchronizer Web



**Events to synchronize**  
*Figure 1: You configure Data Synchronizer filters to determine what data you want synchronized*

Admin console. For example, for the GroupWise connector you can choose from a list of add, modify, move and delete events that it will synchronize. (See [Figure 1](#).)

### > More on Filters

One of the standard filters in [Novell Data Synchronizer](#) provides a circular sync check. As data passes from the connector to the synchronization engine, this filter makes sure that it's not synchronizing the same data over and over again in a circular fashion. This circular check filter prevents the creation of duplicate information.

The default filters that come with a connector are typically sufficient to enable the translation of data from an application-specific format into an application-neutral format and vice versa, as well as apply other basic filter operations. However, since these filters are made up of XSLT files, if you're familiar with XSLT you can manually configure them if necessary. (See [Figure 2](#).)

While you can define additional filter functions using the connector's custom namespace, the following are standard functions provided in the Novell Data Synchronizer namespace:

```
def logMessage(self, context, message, level='warning'): log a message

def getSettingsForTarget(self, context, targetDn): returns a user's settings

def getConnectorSettings(self, context): returns a connector's settings

def getEngineSettings(self, context): returns the sync engine's settings

def getConfigEngineSettings(self, context): returns config engine's settings

def getMatchingUserDN(self, context, sourceName): returns the dn of a user with a
unique cn provided

def dnExists(self, context, dn): checks if the dn is a valid dn in the tree

def isTargetEnabled(self, context, dn): checks if the user is in the user table
for the connector

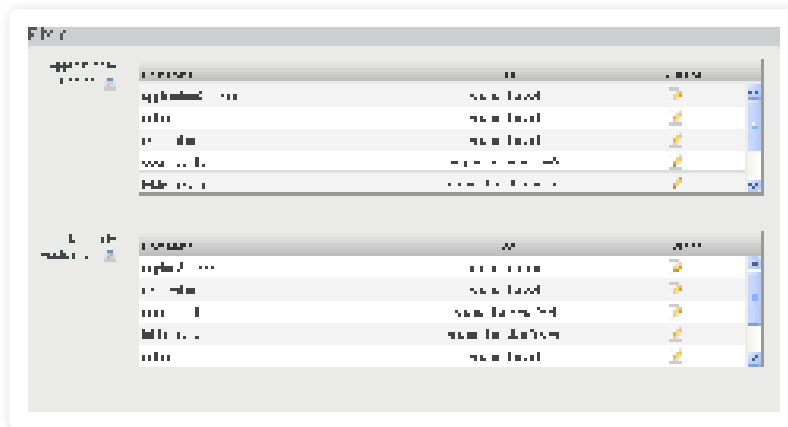
def base64Encode(self, context, message):

def base64Decode(self, context, encodedMessage, returnAsNodeset=False):

def stripHTML(self, context, text, unescapeFirst=False): removes html from the
text

def itemIDToObjectID(self, itemID):

def objectIDToItemID(self, objectID):
```



### Filter management

**Figure 2:** If desired, you can manually configure the filters' XSLT files.

### > Other Novell Data Synchronizer Components

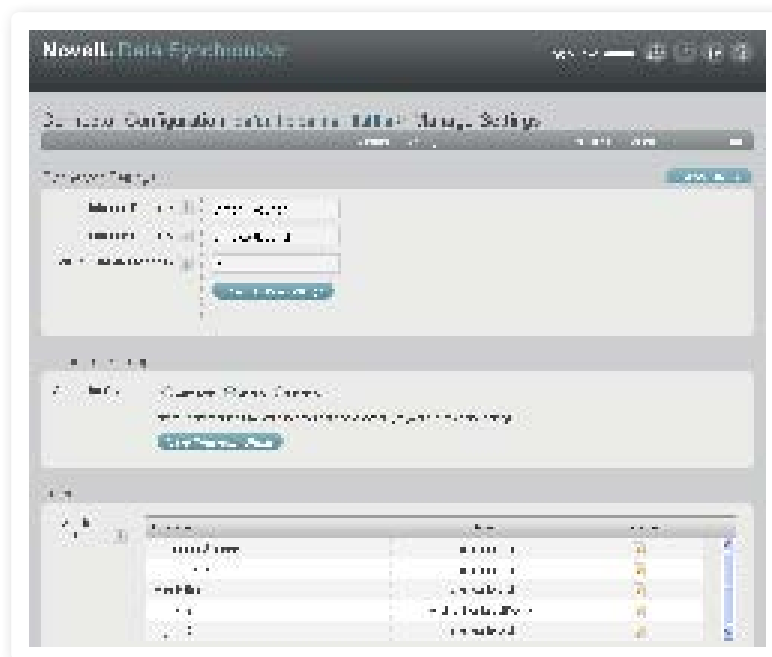
In addition to the synchronization engine, connectors, channels and filters, the following three components play a critical role in the operation of the Novell Data Synchronizer system:

- **Web Administration Service** provides a Web-based interface called Synchronizer Web Admin to facilitate administration and management of the Data Synchronizer system. It allows you to add and remove connectors, as well as add or remove users for specific connectors and configure their individual synchronization settings. Users can also use the Synchronizer Web Admin to configure and control synchronization settings for each of their connected applications and devices.
- **Synchronizer Configuration Engine** provides communication between the Synchronizer Web Admin and the synchronization engine. The configuration engine passes configuration information from the synchronization engine to the Synchronizer Web Admin for viewing. It also passes configuration changes back to the synchronization engine for implementation.
- **Connector Manager** provides communication between the synchronization engine and connectors.

Novell Data Synchronizer also includes LDAP integration. While not required, Novell Data Synchronizer can leverage an LDAP directory to further facilitate user and group management.

---

**In addition to determining what data types are synchronized, filters have responsibility for managing the synchronization of changes that occur to data items within a connected application, system or device.**



**Flatfile connector configuration**

**Figure 3:** The flat file connector can take item and event data cached in the synchronization engine, and translate it into an XML format, and then push it to a specified directory.

**> Developing Connectors**

In addition to the connectors that Novell provides, Novell gives you a number of tools to facilitate the development of additional connectors. The first of these tools is the software development kit for Novell Data Synchronizer, which guides you in how to take advantage of the product's open API.

Novell also provides a generic SOAP connector, which accesses Novell Data Synchronizer using standard SOAP calls. Primarily intended to be used as an example connector, the SOAP connector provides a guide or reference point for someone who wants to develop a connector. You can use the connector as a model to see how the synchronization engine transforms data coming out of the GroupWise system.

As an additional connector development tool, Novell also provides a flat file connector. The flat file connector can take item and event data cached in the synchronization engine, translate it into an XML format administrators can read, and push it to a specified directory or folder on the network. Once placed in the specified folder, it can be consumed by other applications or manually accessed by users. (See [Figure 3](#).)

---

**The synchronization engine, connectors, filters, channels, open development environment and other architectural components of Novell Data Synchronizer all combine to give you the real-time synchronization that your users need among their diverse applications and mobile devices.**

The flat file connector also has the ability to periodically poll a specified folder for newly stored data. When it identifies new data, it can grab that data and pass it back to the synchronization engine for consumption by connected applications or devices.

While it operates on a fairly simple and straightforward concept, its inherent flexibility and open nature make the flat file connector quite powerful in helping you extend the capabilities of Novell Data Synchronizer in a variety of ways. In essence, it gives you a semi-technical programming interface that can help you customize and extend the capabilities of Novell Data Synchronizer.

Leveraging the connector, you could create a custom application that when a certain event happens, it generates information in an XML file format and then stores it in a specified folder. Once in the folder, the flat file connector can grab it, transform it and store it in the synchronization engine where it can be used by other connected applications.

Another possible use would be to use the flat file connector to retrieve and forward reports generated by a certain application. The application could generate the report as an XML file and store it in a specified folder. As the flat file connector polls that folder, it would discover the new report, transform it into an application-neutral format, and store it in the synchronization engine database where it can be consumed by other connected applications or sent to a specific user.

The flat file connector could also be used to pre-populate a new user's GroupWise, SharePoint, or SugarCRM system with standard information. To accomplish this, you might manually create a file in the proper XML format with the desired information and then store it in a specified directory. The flat file connector can then grab, transform and store that data in the synchronization engine cache. When a new account is created for a user, its application-specific connector can then automatically pull that data from the synchronization engine and pre-populate the user's account accordingly.

### > **Syncing It All Together**

The synchronization engine, connectors, filters, channels, open development environment and other architectural components of Novell Data Synchronizer all combine to give you the real-time synchronization that your users need among their diverse applications and mobile devices. They work together to help you improve your organization's collaboration efforts through seamless and real-time synchronization of e-mail, calendar, contact, event and other collaboration data residing in your business and collaboration systems. Whether in the office, at home or on the road, Novell Data Synchronizer can help ensure that your users always have access to the right information in the right place at the right time. To learn more visit [www.novell.com/data-synchronizer](http://www.novell.com/data-synchronizer).

---

#### **Additional Resources** // Novell Connection Magazine

#### **Learn More About Novell Data Synchronizer**

- [Novell Data Synchronizer](http://www.novell.com/products/data-synchronizer) (www.novell.com/products/data-synchronizer)

# Population Automation

## A Successful BSM Depends on an Accurate CMDB

By Bill Tobey

---

**If you're implementing [Business Service Management \(BSM\)](#) for any of the usual reasons—to increase the reliability and availability of your critical business services, to improve your service delivery performance, or to reduce the costs of service delivery—be aware that your success in achieving any of these objectives will be highly dependent on the quality and timeliness of the data in your Configuration Management Database (CMDB).**

Today's CMDB is operational by definition. When an outage occurs, you can't do fast root-cause and impact analysis unless your CMDB can access and supply near real-time information about all services and configuration items (CIs) under management, the relationships between them, and their current service states. The CMDB is a special-case database that must provide four critical functions creating the Configuration Management System (CMS):

- Federation – Direct linkage to multiple data sources
- Reconciliation – Automated data coalescence and matching across sources to eliminate duplication
- Visualization – Automatic detection and graphic representation of the hierarchical relationships between physical and logical CIs
- Synchronization – The ability to immediately capture changes in the physical and logical infrastructure, preserving and providing a single version of the truth across all integrated systems

One of the key challenges in creating and maintaining a CMDB is the problem of populating it from the various system, asset and network management solutions where information about the IT environment originates. Manual processes may offer a high level of initial accuracy, but in large environments, their labor intensity virtually guarantees decreasing accuracy over time.

### > **SCM: Population Propagation, the Automated Way**

But if you're implementing a Novell CMDB, you've got an integrated population automation solution sitting right under your right mouse button, ready to launch. The Service Configuration Manager functions as an intermediary between the CMDB and the adapters that integrate it with management applications and configuration information repositories, providing the Configuration Management System (CMS) automation. The Service Configuration Manager feature dynamically generates new element hierarchies from multiple sources. It automates CMDB synchronization to optimize data quality and reliably capture change in the IT infrastructure in near real time.

The Configuration Management System provides:

- Integration into tools that detects dependencies, characterizes normal baselines, and identifies both scheduled and unscheduled changes
- Integrated mapping of element relationships across an enterprise IT environment
- Integration of asset, configuration and change data from multiple sources into a global CMDB
- Automatic creation and dynamic maintenance of Business Service Views that eliminate manual modeling, mapping and maintenance.

CMS not only automates many of the most time- and labor-intensive aspects of configuration data management, it lets you create and store customized business rules, in the form of JavaScripts, to govern key processes. Let's look at an example.

### > **From Many Sources, One CMDB**

The fundamental use case in CMDB implementation is populating the configuration data store from a large number of diverse management systems and data repositories across the environment, and potentially from other siloed CMDBs. Let's assume that you're starting such a project using the Novell Configuration Management System. You have the Novell engine up and running on a server, along with adapters for each of the management systems you'll need to integrate with. You also have the Service Configuration Manager administrative client running on a desktop or notebook. You want to begin federating basic information about configuration items—services and all the infrastructure elements that participate in their delivery—in the CMDB.

- Begin in the Java console by drilling down into the CMS service model. Select the CI container inside of a community—servers for instance—and right click. From the list of options that is presented, choose Service Configuration Manager (SCM). Then choose to create an SCM job.
- Service Configuration Manager will now display a list of the available source system adapters. Since we're seeking information on servers, select all the adapters for target systems that are likely to contain identity, role, configuration or service state information on server systems.
- Now it's time to decide what information we'll select from each of the available sources, and how we'll load that information into the CMDB. Service Configuration Manager provides a checkbox to populate the CI Attributes from the data sources, but in most cases, unless you have an exact naming match between the source attribute names and the destination CI attribute names, you will need to set up a script to do the actual processing. Regardless, the java script approach provides additional control over how CI Attributes are populated such as pre-processing to do validity checks.

To manage the selection, linking and reconciliation of CI attributes, SCM gives us the option of using a custom JavaScript, like the following:

```
// This script is called for different types of matches. It checks to see what elements exist
// in our source systems, and that each element has been accurately created in the CMDB.
// It begins with three conditions for execution.
if( this.source && element )
{
  // Define attributes (properties) that are on the source system
  if( !state.sourceAttrs )
    state.sourceAttrs = [ 'filename', 'name', 'ip_address' ]

  // Define the attributes (in matching order) that you want the source to populate
  if( ! state.destinationAttrs )
    state.destinationAttrs = [ 'Filename', 'DB Name', 'TCP/IP Address' ]

  //Using state variables above reduces memory utilization on the server.
  //Since this script is run over and over again for each matched element,
  //it reduces run time by not redefining these variables every time.
  //Use a for loop to traverse through the identified source attributes
  for( var i = 0; i < state.sourceAttrs.length; ++i )
  {
    // attempt to pull source attribute
    var value = source[ state.sourceAttrs[i] ]

    // if it exists
    if( value )
      element[ state.destinationAttrs[i] ] = value
    // then on the line above, we then set the element
    // attribute to the value of the source attribute
    // in turn, this is then persisted to the configStore (database)
  }
}
```

Additional information on scripting for SCM control can be found in the FormulaScript Guide that is included in the CMDB installation documentation.

### **> A Script for Accurate, Efficient, Cost-Effective Service Management**

The built-in automation capabilities of the Novell Configuration Management System combined with Service Configuration Manager's support for customized JavaScript control take many of the complexity, labor intensity and data quality challenges out of creating and maintaining a configuration management database. They're just one instance of the extensive automation capabilities that help Novell BSM solutions deliver faster time-to-value, lower implementation and maintenance costs, greater accuracy, improved business value and operational performance from all services under management. To learn more visit [www.novell.com/bsm](http://www.novell.com/bsm).

# From Microsoft Windows XP to Windows 7

## Six Steps to Success

By, Richard Hanley

---

**Migrating Microsoft Windows XP workstations to Windows 7 is a huge job. The pressure to get the job done is intensifying because Microsoft will stop supporting Windows XP SP2 after July 13, 2010 and Extended Support for Windows XP SP3 will end in April of 2014.**

An in-place migration from Windows XP to Windows 7 isn't possible because the two operating systems are based on different architectures. Instead, Windows 7 requires a completely new install. It's not a trivial undertaking, especially if you have to migrate hundreds, even thousands of workstations.

IT organizations that are managing their Windows XP workstations with Novell ZENworks Configuration Management have an advantage: Novell has developed a straightforward, six-step methodology for migrating ZENworks Configuration Management-managed Windows XP workstations to ZENworks Configuration Management-managed Windows 7 workstations. (ZENworks Desktop Management users can also take advantage of the methodology but they must first move up to ZENworks Configuration Management.)

The methodology employs both Novell and Microsoft tools that automate many of the migration processes, minimizing risk and reducing effort and cost. The methodology retains the individual personalities of the user workstations, so users see a smooth transition. (A workstation's personality generally consists of system settings that govern the look and feel of the desktop, configuration data for the user's applications, personal files and Internet browser settings and bookmarks.)

The methodology enables you to use the same ZENworks Configuration Management tools with Windows 7 that you are currently using with Windows XP. The ability to manage a mix of Windows XP and Windows 7 workstations means you can migrate the workstations incrementally.

Novell offers a one-day, instructor-led (online or in-classroom) training course to guide you through the migration process. The course includes lab exercises that provide hands-on experience in using the tools and applying the methodology.

### > The Tools

The methodology uses both Novell ZENworks Configuration Management tools and Microsoft tools. Here's a brief description of each tool used.

#### ZENworks Configuration Management Tools

- ZENworks Configuration Management Personality Migration captures and records the personalities of the Windows XP workstations so you can restore those personalities as part of the Windows 7 migration. This tool generates an XML document that the Personality Migration Engine uses when gathering or applying personality.
- ZENworks Configuration Management Inventory/Reporting Feature enables you to assess

your workstation inventory to determine which Windows XP workstations are capable of accommodating Windows 7.

- ZENworks Configuration Management Bundle Management provides the ability to gather, package, deliver and install content on remote workstations. You can use this feature for such purposes as driving the capture and restoration of workstation personality and installing applications.
- ZENworks Configuration Management Image Explorer is used to build add-on images. These images are so named because, unlike base images, they are deployed without deleting and recreating existing partitions. The contents of add-ons are simply added to the contents of existing partitions. An example of an add-on image is one used to deliver the ZENworks Adaptive Agent to the target workstations.

#### Microsoft Tools

- Windows Automated Installation Kit (WAIK) comprises a set of GUI and command-line tools with associated documentation. Network administrators use it to create, configure and modify Windows 7 operating system images, and to automate the installation of Windows 7 operating system images on target machines. The Windows System Image Manager (WSIM) component exposes all configurable settings of an image file and enables you to write your selected settings to an XML Answer File. Both Sysprep and Windows Setup use the Answer File as input to their processes.
- Sysprep is installed as part of all versions of Windows 7. (You must run Sysprep before deploying a Windows 7 image.) Sysprep prepares a Windows 7 workstation for imaging and removes the machine's "uniqueness." It's generally used as a command-line utility but you can also use it with a GUI.

### > The Methodology

The six-step methodology from Novell wraps around Microsoft Image-Based Setup (IBS) technology that applies an entire Windows 7 image in Windows Imaging (WIM) format to the workstation's drive. You can deploy the image that Windows Setup applies to the drive using either traditional ZMG-based imaging or ZENworks Configuration Management 3rd Party Imaging Support.

#### Step 1 - Identify Windows 7-Ready Workstations

This first step is an assessment of the inventory of Windows XP workstations to determine which ones are capable of migrating to Windows 7. You can use either the ZENworks Configuration Management Inventory/Reporting Features or Microsoft's Windows Upgrade Advisor for this assessment.

#### Step 2 - Configure the Technician's Computer

In this step, you build and configure the Technician's Computer that you will use to create the control files and bundles needed to perform the migration. You begin by installing the operating system. (You can use either Windows XP SP3 or Windows 7. Novell recommends Windows 7.) You then install WAIK and ZENworks Configuration Management Image Explorer, and install and configure ZENworks Configuration Management Personality Migration.

You will use the Technician's Computer for a variety of purposes, including:

- Creating a Sysprep Answer File to configure how the Windows 7 Standard Operating System (SOS) will be installed on the target workstations
- Determining what workstation settings and data are to be backed up and restored as part of the workstation personality migration
- Building ZENworks Configuration Management Application Bundles for installing ZENworks Configuration Management Personality Migration, capturing and applying workstation personality, and deploying images to the workstations
- Building add-on images using ZENworks Configuration Management Image Explorer

#### [Step 3 - Backup Operating System Settings and User Data Files](#)

This step employs ZENworks Configuration Management Personality Migration and ZENworks Configuration Management Bundle Management to save the operating system settings and user data files of all Windows XP workstations so you can restore them to the Windows 7 workstations.

#### [Step 4 - Build and Configure the Reference Computer](#)

At this point, you build a Reference Computer from which you create the Windows 7 Standard Operating System Image that you will apply to the target workstations. You begin by installing your selected version of Windows 7. The configuration of the Reference Computer is determined by the Sysprep Answer File that you created on the Technician's Computer in Step 2. You specify the name of the Sysprep Answer File you created when you run the Sysprep utility on the Reference Computer. After running Sysprep, you then make an image of the Reference Computer using ZENworks Configuration Management ZMG-based or ZENworks Configuration Management 3rd Party Imaging.

Typically, in addition to the Windows 7 SOS base, you will deliver content, such as the ZENworks Adaptive Agent and the Novell Client for Windows 7, to the target machines during the imaging operation. To do so, you create nondestructive add-on images using ZENworks Configuration Management Image Explorer. You also configure a default profile that will be assigned to all local users created on the target workstations.

#### [Step 5 - Deploy Windows 7 SOS Image to Target Machines](#)

In this step, you use ZENworks Configuration Management to deploy the Windows 7 SOS base image created from the Reference Computer, plus any add-on images you created, to the target machines. These add-on images include the ZENworks Configuration Management Adaptive Agent to make the target machines manageable from ZENworks, and may also include other standard line-of-business applications.

#### [Step 6 - Restore User Settings and Data Files](#)

In this final step, you use ZENworks Configuration Management Personality Migration and ZENworks Configuration Management Bundle Management to restore the user data and settings that you saved in Step 2 to the migrated machines. After the target machines have been imaged, they will automatically reboot. Windows Setup will then configure Windows 7 according to the settings in your Sysprep XML Answer File that will also include the installation of the ZENworks Adaptive Agent. Once the agent is installed, the target machines will again reboot. As each target machine reboots, a device object for that machine will appear in the ZENworks Control Center making the machine fully manageable by ZENworks Configuration Management.

## > The Training

The Novell training course presents the generalized methodology, which you can easily customize to meet your organization's specific requirements. There are two recommended prerequisites for attendees. First, they should be experienced Windows users or support people. Second, they should have a basic understanding of ZENworks Configuration Management, including an understanding of the feature set and experience using the ZENworks Control Center.

### Major Topics

The major topics covered include:

- Understanding the new Windows 7 Image-Based Setup (IBS) Technology — Includes an understanding of how IBS differs from Windows XP Setup and an understanding of Windows Imaging Format (WIM)
- Introduction to the Novell Migration Methodology — Includes an overview of the basic processes of the methodology and an introduction to the tools used
- Understanding the Microsoft tool set — Includes Windows System Image Manager (WSIM) and Sysprep
- Understanding ZENworks Personality Migration — Includes a definition of workstation personality, and an understanding of the components of ZENworks Personality Migration and ZENworks Personality Migration Template Builder
- Building a Windows 7 Standard Operating System (SOS) image — Includes selecting the image format, building the Sysprep Answer File that is used to configure the Sysprep process that you will run on your Reference Computer, and running Sysprep
- Creating add-on images to deploy with the Windows 7 SOS Image — Includes installing and using ZENworks Configuration Management Image Explorer, and building add-ons for the ZENworks Adaptive Agent and other third-party applications
- Deploying the SOS base image and add-ons — Includes creating the Image Bundle to deploy your Windows 7 SOS base image and add-on images, configuring the imaging hardware rule to trigger imaging operations, and building a ZENworks Directive Bundle to restore the personalities from the users' Windows XP machines to their new Windows 7 machines

### Lab Sessions

Lab sessions include:

- Session 1. Configuring the Technician's Computer
- Session 2. Configuring ZENworks Configuration Management Personality Migration
- Session 3. Building the Sysprep XML Answer File
- Session 4. Configuring and imaging the Reference Computer
- Session 5. Creating add-on images using ZENworks Configuration Management Image Explorer
- Session 6. Deploying the new Windows 7 SOS environment with add-ons

**> Take the Next Step**

Find out more about how the Novell methodology and training class can help you lower the cost, effort and risk involved with Windows XP to Windows 7 migration.

Visit <http://www.novell.com/training/online/>.

---

**Online Resources** // Novell Connection Magazine

**Learn More about Novell Training and Novell ZENworks**

- [Novell Training](http://www.novell.com/training/online/) (http://www.novell.com/training/online/)
- [Novell ZENworks Configuration Management](http://www.novell.com/products/zenworks/configurationmanagement/) (http://www.novell.com/products/zenworks/configurationmanagement/)
- [Movers Toolbox](http://www.novell.com/connectionmagazine/2009/12/movers_toolbox.html) (http://www.novell.com/connectionmagazine/2009/12/movers\_toolbox.html)


# Raise Your Workload IQ

Start Making All of Your Workloads More Intelligent

by Todd Swensen

Remember when “intelligent workload management” meant postponing a conference call so you could head home early on a Friday afternoon—and “cloud computing” meant playing Solitaire on your laptop on the flight from Chicago to New York? It goes to show that enterprise computing has undergone some major shifts over the past few years, driven by the rise of virtualization and the development of new outsourced services models. And as with any rapid shift in the IT landscape, these changes have introduced some tough questions and—many would argue—the potential for a lot of new IT complexity and uncertainty.

Fortunately, data center virtualization, intelligent workload management and \*aaS—which are all obviously interconnected—are beginning to come into sharper focus. As these relatively new markets mature, vendors and IT departments are working through the issues and developing practical ways to manage and apply them in ways that make sense. In other words, the industry is learning how to make trends like cloud computing look more like practical, serviceable extensions of real-world IT infrastructures and less like a replay of the dotcom boom of the late 90’s. As a result, adoption of these technologies is beginning to accelerate at a fairly impressive rate. (See [Figure 1.](#))



## Physical, virtual and cloud utilization

*Figure 1: The rapid growth of virtual and cloud computing are making intelligent workload management an unavoidable necessity.*

So what are some of the key factors that are steadily transforming intelligent workload management and cloud computing from interesting new trends into practical, mainstream enterprise solutions? There are a few key concepts and practices that most experts agree will fuel the widespread adoption of these technologies over the next few years.

### > **The Evolution of the Workload**

There's certainly nothing new about the concept of a workload. We all know it's simply an integrated software stack that includes just enough operating system, middleware and some kind of application. But to work well in virtualized or cloud environments, these workloads need to evolve to become more portable and platform agnostic, so they can be combined into business services and deployed quickly into any kind of physical, virtual or cloud infrastructure. For example, you should be able to quickly deliver a business service that combines a data base workload running on physical hardware in a legacy data center, an application server running in a virtualized private cloud, and a presentation and graphics server running in a public cloud. And you should be able to move these workloads quickly and automatically across these different environments whenever it makes sense. To meet these new portability requirements, many enterprises are turning to software appliances as a fast and practical way to package, configure and deploy, or in other words build self-contained workloads into physical, virtual and cloud environments. This approach makes sense, but building more portable workloads is only part of the equation. You also have to find effective ways to manage, secure and measure those workloads as they move across these different environments.

### > **Gaining Flexibility without Losing Control**

---

**Info** // Novell Connection Magazine

#### WHAT MAKES A WORKLOAD INTELLIGENT?

What are the biggest differences between the traditional workloads of the past and the intelligent workloads of the future? Intelligent workloads are:

- **Policy-driven**, which means they can regulate and manage themselves, recognize when they are at capacity, and automatically find additional capacity—all based on pre-defined policies.
  - **Secure**, which means they include security controls and real-time monitoring and alerting that move with them between environments.
  - **Compliant**, which means they understand and follow security protocols and processing requirements and provide built-in log management and compliance reporting.
- 

The ability to instantly and automatically move workloads among different physical, virtual and cloud resources depending on the situation has immense appeal and offers obvious benefits for lowering costs and improving service levels. But that flexibility simply can't come at the expense of operational control or increased risk. To take advantage of portable workloads, enterprises have to find ways to extend all of their essential operational and risk control mechanisms across all their different physical, virtual, and internal and external cloud environments.

**Info** // Novell Connection Magazine

## NOVELL WORKLOADIQ SOLUTIONS

Novell can increase your WorkloadIQ with solutions that help you build, secure, manage and measure workloads across all your physical, virtual and cloud environments:

- **Build** with SUSE Linux Enterprise Servers and the SUSE Appliance Program.
  - **Secure** with identity and access management, compliance management, and security management for intelligent workloads.
  - **Manage** with virtualization, workload, and endpoint management solutions for intelligent workloads.
  - **Measure** with complete business service management solutions for intelligent workloads.
- 

This isn't easy, because in most cases physical, virtual and external cloud environments have separate and siloed governance and compliance, business service management, and IT service management mechanisms in place. Enterprises are justifiably reluctant to move workloads from their secure, compliant and meticulously managed legacy data centers to an external cloud environment where they have to trust a third-party, off-premise vendor to provide those same services. So what's the answer? In very simple terms, you have to break down the boundaries, so your workloads can access the same management, security, monitoring and compliance mechanisms across all your physical, virtual and cloud computing environments.

### > **Adding Identity to Every Workload**

Of course, that's not as easy as it sounds, especially when many of those mechanisms are ultimately controlled by third-party cloud providers. In a world where workloads need to move freely and securely among physical, virtual and cloud environments, it no longer makes sense to develop complex rules engines and workflows that can work across every possible type of environment and infrastructure. Instead, the Novell approach—called [WorkloadIQ](#), embeds security, management, monitoring and compliance controls inside the workloads themselves. In a way, this approach does for workloads what identity has done for individual users, and it's the key component that makes workloads intelligent. By embedding identity-like attributes into workloads, you make it possible for them to act according to business policies, automatically find alternative computing capacity to optimize performance, adhere to established security controls in any environment and much more. You also embed a kind of workload RFID tag that can provide real-time tracking, monitoring and alerting for every workload—even if it's running in an off-premise cloud environment. Ultimately, this identity-based approach to workload management is the key to breaking down the silos between environments, making workloads totally portable and imbuing them with all the attributes and capabilities they need to stay compliant, safe and well managed in any environment.

**> WorkloadIQ: Build, Secure, Manage, Measure**

The effort by Novell to make workloads more portable and intelligent revolve around a new WorkloadIQ strategy—backed by a family of products and solutions designed specifically to build, secure, manage and measure the kinds of flexible, intelligent workloads you need to take advantage of physical, virtual and cloud computing. Together, these technologies can help you build intelligent, policy-driven, identity-aware workloads that will enable you to take full advantage of everything cloud computing and virtual technology can offer your business.

For more information on how Novell WorkloadIQ can help you build, secure, manage and measure intelligent workloads, visit [www.novell.com/workloadiq](http://www.novell.com/workloadiq) or visit our blog at [www.workloadiq.com](http://www.workloadiq.com).

---

**Online Resources** // Novell Connection Magazine

**Learn More about WorkloadIQ**

- [WorkloadIQ](http://www.novell.com/workloadiq) (www.novell.com/workloadiq)
- [WorkloadIQ Blog](http://www.workloadiq.com/) (http://www.workloadiq.com/)

# Does Apple's iPad Really Mean Business?

Love It or Hate It, You Should Probably Start Paying Attention to Apple's iPad

by Todd Swensen and Eric Harper

---

**What was your gut reaction when Steve Jobs first unveiled Apple's new iPad tablet computer on January 27th? Fawning admiration? Healthy skepticism? Open hostility? All of these reactions were well represented and expressed in a frenzy of articles, blog posts, tweets and coffee shop conversations. When the iPad finally started showing up in retail stores in early April, millions of people already wanted one (although most couldn't have told you exactly why). Since then, surprisingly robust iPad sales—together with unprecedented interest in developing apps for the iPad—seem to indicate that something big is happening. So whether you consider the iPad another overhyped consumer fad built for people who like to show off at Starbucks, a legitimate technology game changer, or something in between, it's hard to argue with success. And that probably means it's time to at least start asking questions about what the iPad might mean for your business.**

## > Gauging the iPad's Potential

Before the iPad shipped, technology experts and reviewers joined millions of consumers in wondering exactly what this heavily hyped new device was and how people and businesses might actually use it. Would it suffer from the same limitations (and limited success) as other tablet PCs? Would it offer something unique—an experience you couldn't get from a less expensive smartphone, netbook or other portable device? Would it have anything compelling to offer businesses? Now that the iPad is actually available, experts, reviewers and users have started chiming in on some of these questions. Wall Street Journal tech reviewer Walter Mossberg said, "After spending hours and hours with it, I believe this beautiful new touch-screen device from Apple has the potential to change portable computing profoundly, and to challenge the primacy of the laptop. It could even help, eventually, to propel the finger-driven, multitouch user interface ahead of the mouse-driven interface that has prevailed for decades." (WSJ article March 31, 2010)

---

Info // Novell Connection Magazine

## NOVELL MONOTOUCH 2.0—IPAD APPS FOR .NET DEVELOPERS

Microsoft is not currently developing any iPad apps, but thanks to Novell, the same does not hold true for Microsoft .NET developers. Novell MonoTouch 2.0 provides .NET developers with the tools they need to:

- Develop applications for the Apple iPad, iPhone and iPod Touch using Microsoft .NET languages
  - Extend iPad development to .NET developers around the world
  - Create apps that turn the iPad into an efficient, productive business tool
-

On the business side, Salesforce.com chairman and CEO Marc Benioff believes the iPad has important implications for the next generation of cloud computing. In a recent guest post on TechCrunch, he said, "What's most exciting is that this fundamental transformation—cloud + social + iPad—will inspire a new generation of wildly innovative new apps that will change entire industries... This will result in a new generation that looks more like Facebook on the iPad than Yahoo on the PC." (<http://techcrunch.com/2010/03/29/ipad-cloud-2/>)

There are also some legitimate iPad concerns and criticisms to counter the enthusiasm—especially in a business context. Many worry about Apple's relatively closed application development and distribution model, its restrictive DRM controls, and the company's overall propensity for control over openness. And of course, the iPad introduces the same security and management challenges as any new, untested portable device. These concerns are causing many businesses to take their time with the iPad. But given its fairly remarkable early success, it's also becoming clear that the iPad is not something businesses can afford to ignore.

### > **Putting the iPad to Work**

Some of the industry's most prominent enterprise software vendors are certainly paying attention. PeopleSoft, SAP, Oracle, Novell and many others are actively working on developer kits and enterprise iPad applications. In fact, Oracle already has eight iPad applications available through Apple's App Store. Many of these early CRM and database apps are nearly identical to their existing iPhone counterparts, but developers are building specific business apps that take advantage of the iPad's beefy processor and other unique capabilities. This level of activity from major enterprise application providers is notable, given that the iPad has only been available for a few short weeks.

---

**Info** // Novell Connection Magazine

#### IPAD IN THE ENTERPRISE

- Oracle Announces Siebel CRM Support for the iPad (link to: <http://www.oracle.com/us/corporate/press/068482>)
- Salesforce CEO Marc Benioff on the iPad and Cloud 2.0 (link to: <http://techcrunch.com/2010/03/29/ipad-cloud-2/>)
- IBM launches Lotus software for the iPhone platform (link to: <http://www.forbes.com/2010/02/11/ipad-iphone-apple-technology-cio-network-ibm.html>)

### > **The Microsoft Exception**

The most significant exception to this trend is Microsoft, which to this point has chosen not to develop any applications—including Microsoft Office—for the iPad. Stephen Elop, the president of Microsoft's business division, recently told Bloomberg that they were taking a "wait and see" approach to the iPad. "We never say never," said Elop, "but we have no current plans to develop a version of Office for the Apple iPad."

At one level, this makes sense. Microsoft doesn't typically like potential game changers, because they prefer the game (which they currently dominate) just the way it is. It will be interesting to see how this strategy plays out—especially given that Apple's Pages word processor is already available for iPad. Google, which is directly targeting Microsoft Office dominance with Google Apps, is also actively working to optimize its offerings for the iPad (despite intriguing rumors of an upcoming "[Google iPad killer](http://gizmodo.com/5514989/google-preparing-ipad-rival)"). (<http://gizmodo.com/5514989/google-preparing-ipad-rival>)

---

**Info** // Novell Connection Magazine  
**IS AN "IPAD KILLER" ON THE WAY?**

According to a recent [New York Times article](http://www.nytimes.com/2010/04/12/technology/12slate.html) (<http://www.nytimes.com/2010/04/12/technology/12slate.html>), Google, HP, Microsoft, Nokia, and others are all working on new slate computers that will compete directly with the iPad. The upcoming Google and HP devices will both run Google's Android OS, and Google is apparently working quietly with publishers to explore distribution options for books, magazines and other content. [Link to: <http://www.nytimes.com/2010/04/12/technology/12slate.html>]

---

### > **Off and Running**

When the iPad was first announced, a number of industry watchers and reviewers openly wondered why Apple chose to put such a powerful 1GHz A4 processor into such a simple device. Clearly, Apple was counting on people (and businesses) wanting to use the iPad for much more than reading books, browsing the Web, and watching movies. That bet appears to be paying off. It has been less than 3 months since the iPad became available in the U.S. People apparently want them more than ever. Only now, they actually know why.

---

**Online Resources** // Novell Connection Magazine  
**Learn about MonoTouch from Novell**

- [MonoTouch](http://monotouch.net/) (<http://monotouch.net/>)