



Case Study: Thomson Reuters



Mark Bluhm
Senior VP and CTO, Shared Services
Thomson Reuters Professional Division

At [Thomson Reuters](#), our mission is to satisfy the information needs of businesses and professionals in a wide range of fields, so information technology is critical to everything we do. The seeds of our current approach to IT were planted more than 10 years ago, when we started experiencing stability challenges with our online legal research service, Westlaw.

At that time—before the dot-com bust—Westlaw was still a mainframe-based legacy platform and we were losing talented software engineers who wanted to work with newer technologies. I was tasked to create a new, open infrastructure for Westlaw, and to do it in such a way that the same infrastructure would be able to also support our other information businesses. This turned out to be a pretty farsighted prescription to create shared infrastructure using standard building blocks.

This simple directive put us on a path of steady IT evolution over the years and, most recently, contributed to the successful release of a completely new, next-generation legal research service, WestlawNext. Our infrastructure allowed us to add support for WestlawNext while avoiding roughly \$65 million in new data-center costs, using 25% less power, and delivering 24/7/365 availability. WestlawNext is able to search 50 times more data (5 billion documents) relative to the previous generation and to return results twice as fast.

In this article, I want to highlight some of the important elements of that infrastructure including building blocks, our core search architecture, and our virtualized front end. NetApp and NetApp Professional Services were invaluable partners in this effort, so I'll also try to give credit where credit is due.

A Shared IT Infrastructure for Search

The key to success for WestlawNext and all Thomson Reuters products is to be able to perform searches on massive amounts of data very quickly and with complete accuracy. If two people perform the same search at the same time, they should get exactly the same results.

Because of enhancements to our search methods in WestlawNext, users can simply ask for what they want in plain English—they no longer have to know how to "construct" a formal query. As a result, a query that generated just one search two or three years ago now results in 40 or more searches on the back end, and our infrastructure is still able to scale to meet this load, which is absolutely amazing. It's gone way beyond the original targets that we had for it. A typical search takes just 2.5 seconds to return data to the client.

The key elements of our infrastructure include:

- Standard building blocks
- A cloud-like search architecture
- Virtualized Web front end
- Replication for disaster recovery

Standard Building Blocks

Our infrastructure consists of pretty standard building blocks. We have between 25,000 and 30,000 x86 servers in our data centers, most with 2- or 4-CPU configurations and backed by NetApp® storage. Our network infrastructure is almost entirely 10-Gigabit Ethernet using Cisco 6500 and Cisco Nexus 5000 and 7000 family switches. We use these building blocks in both the front-end and back-end configurations.

Thomson Reuters Key Metrics
25,000+ Servers
NetApp Storage with Flash Cache
Hundreds of Oracle RAC clusters
Novus search infrastructure built on Linux and serving 30+ applications
VMware to virtualize front-end
Avoided \$65 million in new data center costs

Explore

About Thomson Reuters

[Thomson Reuters](#) is the world's leading source of intelligent information for businesses and professionals. The company combines industry expertise with innovative technology to deliver critical information to leading decision makers in the financial, legal, tax and accounting, scientific, healthcare, science, and media markets, powered by the world's most trusted news organization. Headquartered in New York with major operations in London, England, and Eagan, Minnesota, Thomson Reuters employs some 55,000 people in more than 100 countries and reports revenues of US\$13.1B (2010).

We'd love to hear your comments about this article.

Did you find this article useful?

- Strongly Disagree
- Disagree
- Neutral
- Agree
- Strongly Agree

Reduced power consumption 25%
Searches 50X more data (5 billion documents) in half the time

Figure 1) Notable achievements for WestlawNext and the Thomson Reuters IT transformation

Novus: Cloudlike Infrastructure for Search

Our Novus architecture, [which was patented in 2006](#) is the core of all search operations. The Novus architecture provides a single platform for supporting online services from each of the four Thomson market groups, including WestlawNext, and Checkpoint®, our tax and accounting research system. In all, 30+ applications use the Novus architecture.

The Novus system is a distributed search architecture that uses thousands of SUSE Linux® servers each running our proprietary software. Each search server is responsible for part of the overall content index, which fits in server memory so it can be accessed extremely quickly. When a search is executed, it hits thousands of machines at once. The results are sent back to a controller, which sorts them, aggregates them, ranks them, and sends that information back to the requesting application. By doing it this way, we can get subsecond search performance.

The application then decides whether it wants to pull the documents identified in the search. The content stores aren't actually touched until a document is requested. The content itself is stored using hundreds of Oracle® RAC database clusters, typically with four nodes per cluster. Each cluster holds a subset of the total content.

I know that the term "cloud" means different things to different people, but Novus is designed to deliver the flexibility usually attributed to cloud infrastructure, although the infrastructure was designed before the term cloud was popularized. Any server in the Novus environment can be reallocated in real time to take on a different function. When we architected this we wanted to make sure that if a peak event happened, we could reallocate resources very quickly so that, for instance, what was a database server five minutes ago could now be a search server.

When we do code deploys to Novus, all of the code is deployed to every server for every function. So, all we have to do is change a simple setting and say, "Server A, you're no longer a search server, you're a load server."

If WestlawNext is getting hit hard, we can allocate more resources specifically to it, or to Checkpoint or any other application that needs the resources. Servers don't have to reboot, they simply load the appropriate indexes into memory from NetApp storage and they are ready for their new role. Multiple sets of servers can be assigned to the same set of indexes to increase parallelism to allow Novus to continue to scale.

This dynamic capability also allows us to build redundancy into the environment and it ensures result accuracy. We always have extra, idle servers available. If within just a few milliseconds after sending a request we don't get a result back from a server, we do a couple of fast tests of that server. If it doesn't respond, is slow, or is having some other problem, another server will automatically be assigned to assume that role. It will then load the appropriate index into memory and service the request.

The end result is that a server can fail and the user will still get an accurate result with nothing omitted and only a few seconds' delay. The user doesn't have to reissue the request and the recovery happens automatically without administrator intervention. For the Novus content itself, the use of Oracle RAC provides the redundancy. If a RAC server fails, another node in the cluster performs its function. If a RAC cluster is getting hit hard, we can dynamically add more nodes to accommodate the load.

Virtualized Front End

For everything on the front end—everything outside of Novus—we use a much more typical environment consisting of Web servers and various application servers. In addition to accessing Novus for search, the application tier also accesses a variety of things that are not central to this discussion like security databases, user information, billing databases, MIS data, all the things that a normal application needs.

A large part of the front-end environment has been virtualized with VMware®. Most Web servers and application servers run in virtual machines. VMware gives us the ability to do the same kind of dynamic resource allocation on the front end that we do within Novus. We can fine-tune the number of Web servers and application servers for each application as necessary.

VMware also gives us nonstop operation. VMware HA protects against virtual machine failures, and vMotion™ gives us the ability to do maintenance and other operations without any downtime and without losing any in-flight work, which is something we couldn't do before. Like everyone else, before virtualization, if I had 100 users on a server that needed maintenance I would have to quiesce them and take them offline and make them sign back in—or do something magical programmatically, which was nearly impossible.

With VMware, we can do maintenance as necessary during the middle of the day because we can just move running VMs on to an auxiliary set of servers and then do whatever maintenance we need to do on the original servers.

Disaster Recovery

I've already explained how we provide redundancy within a data center, but I've held off talking about disaster recovery (DR) to keep things simpler. Under normal operation, we always have two data centers running with very similar infrastructure and identical data. If a disaster takes down one running data center, the other running data center can scale up operations to accommodate the additional search load.

We use replication to keep our data centers in sync. We have our own replication mechanisms that we've developed to support replication of our Novus indexes and make sure they are perfectly synchronized. The content stores in our Oracle RAC databases are replicated using Oracle DataGuard.

NetApp Changes the Game

NetApp storage supports the Novus architecture (indexes and Oracle RAC content stores) as well as the front-end VMware environment. All the indexes that get pulled into our Linux servers plus all the content stored in Oracle RAC are kept on NetApp NAS storage accessed via NFS. Novus simply would not work if we couldn't have thousands of servers sharing access to our storage systems at one time with the ability to dynamically change which servers access which storage on the fly. NetApp storage was a real game changer for us when we first implemented it in 2002, and it remains a critical part of our solution today.

In order to support the scaling and performance requirements of WestlawNext we made some recent infrastructure enhancements. We added [Flash Cache](#) to key NetApp systems. Specifically, we started using these on NetApp systems that provide storage for a single Oracle RAC cluster. Such clusters often have low-capacity and high-performance requirements, so the Flash Cache helps us keep performance high without needing to add spindles and waste capacity to get the required performance. We have also started to use Flash Cache on the shared storage systems that provides the indexes and other data to our Linux clients, and we expect it to have a similarly big impact there based on preliminary testing.

As you might expect, we are adding new content all the time and that means reindexing and pushing both the new content and associated indexes out while keeping everything in sync. If a problem occurs and we need to roll back to a previous state, it has to be done as quickly as possible. NetApp SnapRestore® technology is far and away the best solution we've found to accomplish this.

Before we do a content load, we create a Snapshot™ copy. Then, if we need to roll back for some reason, we can simply do a SnapRestore operation to return our storage (in one data center and then the other) to the state it was in before the load started. (In some cases, for databases, logs may need to be replayed.)

We use NetApp deduplication in our VMware environment to eliminate the duplication that comes with having a large number of nearly identical VMs. One division alone has over 9,000 VMware VMs running on NetApp storage, and we've achieved over 160TB of space savings on primary storage through the use of deduplication.

To manage our environment, we use the full complement of NetApp OnCommand™ management products including Operations Manager, Provisioning Manager, Performance Manager, and OnCommand Insight. This gives us a single set of tools that work across all our NetApp storage to simplify management, speed up provisioning, and identify performance issues. OnCommand Insight (formerly known as NetApp SANscreen®) gives us a consolidated view of our entire heterogeneous storage environment in terms of capacity, connectivity, configurations, and performance. It also provides alerts on component failures so that we can resolve issues before redundant components experience a second failure.

Doing More with Less

I mentioned the significant efficiency and scalability benefits we've achieved by implementing WestlawNext and other services using the infrastructure I've described. By sharing infrastructure on the back end, we are able to efficiently meet peak demand for our various applications by allocating resources where they are needed while keeping idle resources to a minimum. Virtualization on the front end has allowed us to reduce server count and other associated infrastructure there. Our overall efforts have so far saved us from building an additional data center. NetApp storage technologies, including Snapshot copies, SnapRestore, Flash Cache, and the full suite of management capabilities, help us optimize storage use and eliminate bottlenecks.

For Thomson Reuters, our overall relationship with NetApp is just as important to our success as NetApp technology. Of all the vendors we work with, NetApp is one of only two we consider a strategic technology partner. Any problems get fixed immediately, and NetApp is always ready to support us with key technology initiatives like WestlawNext. NetApp has worked closely with us to optimize performance and to help us quickly leverage new storage functionality.

» [Got opinions about the Thomson Reuters case study?](#)

Ask questions, exchange ideas, and share your thoughts online in NetApp Communities.



By **Mark Bluhm, Senior VP and CTO, Shared Services, Thomson Reuters Professional Division**

Mark Bluhm is senior vice president and chief technology officer of Shared Services, overseeing data center operations and strategy for the Professional Division of Thomson Reuters.

Mark has over 19 years of experience at the company, starting at then-West in 1991 as a software engineer. Since then, Mark has held several technology leadership roles, including chief architect for then-Thomson Legal & Regulatory. In fact, he is one of the key inventors of the Novus technology and is the prime patent holder for this TRGR-owned enterprise solution. Following the acquisition of Reuters by Thomson in 2008, Mark was involved in the merger of the two companies' infrastructures. Most recently, he was chief technology officer, Client Development Technology, Legal.

Mark received both his bachelor's and master's degrees in mathematics and computer science from the University of South Dakota, and his doctorate in computer science from Washington University.

Tech OnTap
[Subscribe Now](#)

Tech OnTap provides monthly IT insights plus exclusive access to real-world best practices, tips and tools,

behind the scene engineering interviews, demos, peer reviews, and much, much more.

Visit [Tech OnTap in the NetApp Community](#) to subscribe today.



Go further, faster®



[Contact Us](#) | [How To Buy](#) | [Feedback](#) | [Careers](#) | [Subscriptions](#) | [Privacy Policy](#) | © 2011 NetApp