


# Keeping an even keel

What do sailing and data warehouse management have in common? More than you'd think. *by Dan Higgins*

**H**ave you ever been on a sailing cruise? In 1998, I spent two weeks sailing with my parents on their 36-foot Catalina sailboat called Iwalani II. We began near Seattle, WA, and sailed north into the Canadian waters between Vancouver Island and British Columbia. It was a wonderful trip. While I wasn't thinking about data warehousing at the time, as I reflect on my sailing experiences of the past and on that trip in particular, I see some interesting analogies that can be drawn between sailing and data warehousing.

There's an elegance to a well-designed, well-tuned boat under sail. There is a thrill when you feel the boat surge from the force of the wind driving the matched combination of a fast hull and perfectly trimmed sails. Yet the Iwalani II wasn't just fast, she was also stable. When wind gusts and waves would buffet her, she'd take the surge and then settle back down



to business. You could put her on either a starboard or port tack in a steady 14-knot breeze, trim the sails and let go of the wheel, and the boat would be so balanced she would continue on the same tack.

The shared-nothing parallel architecture of the Teradata Database has always had a certain design elegance of its own. Extensive and experienced engineering yields an optimal match of hardware and software specifically for relational database processing and decision support. This balanced system—together with a well-tuned data model and well-tuned workloads—can yield an incredible amount of value. But it's not enough to have a great data warehousing platform by itself: there is still work you need to do.

## Setting your course

Sail cruising isn't the same as sailboat racing or day sailing. When cruising, you intend to get somewhere, to reach a new destination each day along your journey. To chart your course, you must anticipate the weather, understand the timing and level of the tides as well as the depths of the water, and know the capabilities of your boat and crew.

Similarly, to be effective in data warehousing, you must carefully plan each stage of your journey. These plans need to move you toward long-term goals as well as account for tactical needs. You need to anticipate the potential storms caused by urgent business requirements. And you must understand the ebb and flow of user workloads as well as the capabilities of your technology and staff.

You also need to realize that the best course isn't always the straightest path. Not everyone realizes that you can't sail a boat directly into the wind. Attempting to do so would bring the boat to a quick halt. The same is often true with data warehousing.

I've observed a number of data warehousing efforts that failed or struggled simply because they attempted to sail directly into the headwinds of corporate politics, inadequate technology or lack of available source data. I remember a case where the chief data modeler was removed from a project because the person insisted on the implementation of a "perfect" normalized data model when the necessary source data wouldn't be available for at least six months. There was an unwillingness to compromise, prioritize the implementation of certain model changes and allow the model to iteratively evolve.

In sailing, the best path is to tack a few degrees off the wind and then come about and tack the other direction—all the while moving toward your ultimate goal. Similarly,

in data warehousing you often need to make tactical compromises in order to deliver business value while evolving toward strategic goals.

### Trimming the sails

Even a great boat will be slow when the sails aren't set correctly for the force and direction of the wind. There's always a certain amount of tuning required to achieve the optimal speed.

Getting the most out of a data warehouse platform requires a little tinkering as well. Even a modest amount of tuning can yield huge results in performance and cost-efficient resource utilization. And, as with a sailboat, balanced tuning requires knowledge of the platform and the workloads being serviced by that platform.

As you might expect, this knowledge often comes through trial and error. There are many Teradata customers who would gain far more value from their data warehouse platforms if

they would take the time and effort to do some modest tuning.

### Taking the helm

You need a steady hand on the helm of a boat to steer a straight course; oversteering can significantly degrade the boat's forward momentum. Waves and the occasional wind gusts can rock a boat and temporarily change its course. A good boat will usually settle back in line of its own accord. At other times, a capable sailor can take advantage of changes in the wind to ease the boat into a new direction and capitalize on the new opportunity.

The one thing that you can be sure to anticipate with your data warehouse system is that the unanticipated will happen. Workloads will surge and flow. Killer queries will batter the system like a gust of wind. A good database platform will be resilient to these changes and surges. It will make the necessary adjustments in real time to not only absorb those waves, but, in some cases, to also exploit them. A good data warehousing team will manage the system appropriately and not overreact.

My family's boat, like most cruising boats, has an autopilot which puts it on course and maintains that course once the bearings are set. A good data warehousing platform has various forms of an "autopilot" to manage the system in real time. For example, you can use workload management capabilities to define workloads and set the priorities, and then let the resource and workload management facilities keep the system "on course" while you simply monitor what is going on and manage any exceptions.

As any sailor knows, you can't simply set the autopilot and then go to sleep at the helm. I remember an experience in 1985 with Teradata when someone from our operations group came to me excited about how an AMP (node) on our Model 2 DBC1012 had been down for a week and, because of the system's automated recovery, none of the employees had even noticed. Of course, my first thought

## Teradata Active System Management

Although not quite as predictable as the tides, data warehouse system utilization certainly ebbs and flows. Monday morning, end-of-the-month, first-of-the-month, peak user periods, end-of-day updates and loads—these are just a few examples of how and when system utilization can differ significantly. With the proper monitoring tools, however, you can understand these utilization tides, predict them and even factor them into capacity planning and workload management. While you'll be able to anticipate many of these utilization characteristics, you'll also run into cases, like the weather, that occur without warning and you'll need to be able to respond to those as well.

Teradata has always delivered a very stable and balanced data warehousing platform. Its integrated software and hardware architecture eliminates many of the complex system administration and management tasks associated with other data warehousing platforms. But with increasingly complex and demanding data warehousing workloads to manage, the data warehouse staff needs increasingly capable tools. The Teradata Active System Management suite is a best-in-class example of an "autopilot" for data warehousing workload management.

Teradata Active System Management significantly reduces the effort required by database administrators (DBAs), application developers and support teams to define, monitor and manage data warehousing workloads. In much the same way that a sailor sets his course to autopilot, the data warehousing staff can define workloads based on predetermined goals and expected conditions and "set" the system to manage those workloads. All the data warehouse staff needs to do then is sit back and monitor the system for exceptions or anomalies.

The recently introduced product suite is an evolution, integration and enhancement of previously existing Teradata System Management tools, which includes four key components:

- > **Teradata Workload Analyzer** helps with the analysis of varying workloads based on historical usage data from the Teradata Database. This is analogous to studying tide charts, water depths and weather forecasts before charting your course.
- > **Teradata Dynamic Workload Manager** provides for the creation and management of workload definitions. Think of it as loading your courses into the autopilot.
- > **The Regulator** controls and adjusts the assignment of system resources based on the prescribed workload definitions. Think of this as the autopilot in action.
- > **Teradata Manager** provides important monitoring and exception reporting tools. It's important not to fall asleep at the helm.

## Smooth sailing through the storms and tides of data warehousing requires a stable and balanced platform with robust and reliable system management tools.

was, “Where are the operations people? Is no one paying attention?”

The system, which hadn’t had an outage for months, required almost no involvement on the part of the operations group, who had begun to take the smooth sailing for granted. The lesson learned here? Despite the conveniences of

automation, self-management and built-in recovery, you still need to keep your eyes open for the unexpected.

### Sandbars and shallows

When sailing, you want to keep the sailboat’s keel, which extends a few feet into the

water, from burying itself into a sandbar or slamming into rocks. You may be tempted to cut corners and plot your course closer to land, but you risk spending more than a few hours waiting for the tide to lift your boat off the sandbar.

We sometimes find data warehouse systems that attempt to run at 100% utilization—or sometimes more, in the case of the Teradata Database. With linearly scalable technology, we tend to behave as if the system’s capacity is limitless. We keep adding data and workloads until the system simply has no resources left, and then we wonder why performance is degraded. Good workload and system resource management, monitoring tools and best practices can help, but it’s also important to maintain just a bit of headroom in the system for those unanticipated, yet critical, tasks. And this requires regular capacity planning.

Smooth sailing through the storms and tides of data warehousing requires a stable and balanced platform with robust and reliable system management tools. But to realize the full value of your investment, it still falls to the captain and crew to know their destination, plot their course, tune the system for optimal performance and resource utilization, leverage the built-in “autopilots” and monitor the system for the unexpected. When this is done well, you’ll not only cover great distances, but will also enjoy the sail. **T**

*Dan Higgins is director of Teradata Warehouse sales support and co-author of Evolving Through Action. He leads a team of subject matter experts who assist Teradata’s worldwide sales force.*

### Similarities between sailing and data warehousing

SAILING	DATA WAREHOUSING
Chart each leg of your course given your ultimate destination and factor in wind, weather and water conditions and the abilities of your boat and crew.	Plan each stage of your data warehouse evolution given your strategic goals; consider data availability, workloads and the capabilities of your platform and staff.
Tack—make progress into the wind by sailing a zigzag course.	Evolve the data warehouse and deliver business value in the face of challenges by prioritizing and sometimes compromising on incremental improvements.
Trim the sails in order to maximize boat speed.	Tune the platform, physical data model and workloads in order to optimize resource utilization and maximize performance.
Keep a steady hand on the helm—don’t overreact when the boat is nudged off course due to a wave or wind gust. Oversteering kills momentum.	Don’t overreact to the occasional surprise query or unanticipated business requirement.
Autopilot requires that you know your course.	Fully utilize the available tools to understand, define, manage and monitor varying workloads.
Watch the water for obstacles and for indications of wind and weather changes.	Monitor the system for exceptions—even with automated tools and built-in automation.
Monitor your depth gauge to avoid shallows and sand bars.	Utilize capacity measurement and planning to ensure you have adequate resources to meet business needs.
Experienced captain and crew; you don’t begin with open ocean sailing.	You need experience and expertise; start with the basics and grow from there.
A fast, balanced and stable boat.	Use Teradata!