A MAGAZINE FOR AIRLINE EXECUTIVES 2011 Issue No. 2

ascend

Taking your airline to new heights

The Jewel

A Conversation With … Muhammad Ali Albakri, Chief Information Officer, Saudi Arabian Airlines,

Pg. 24

Continuous Availibility

Talking Technology With … Dolly Wagner-Wilkins

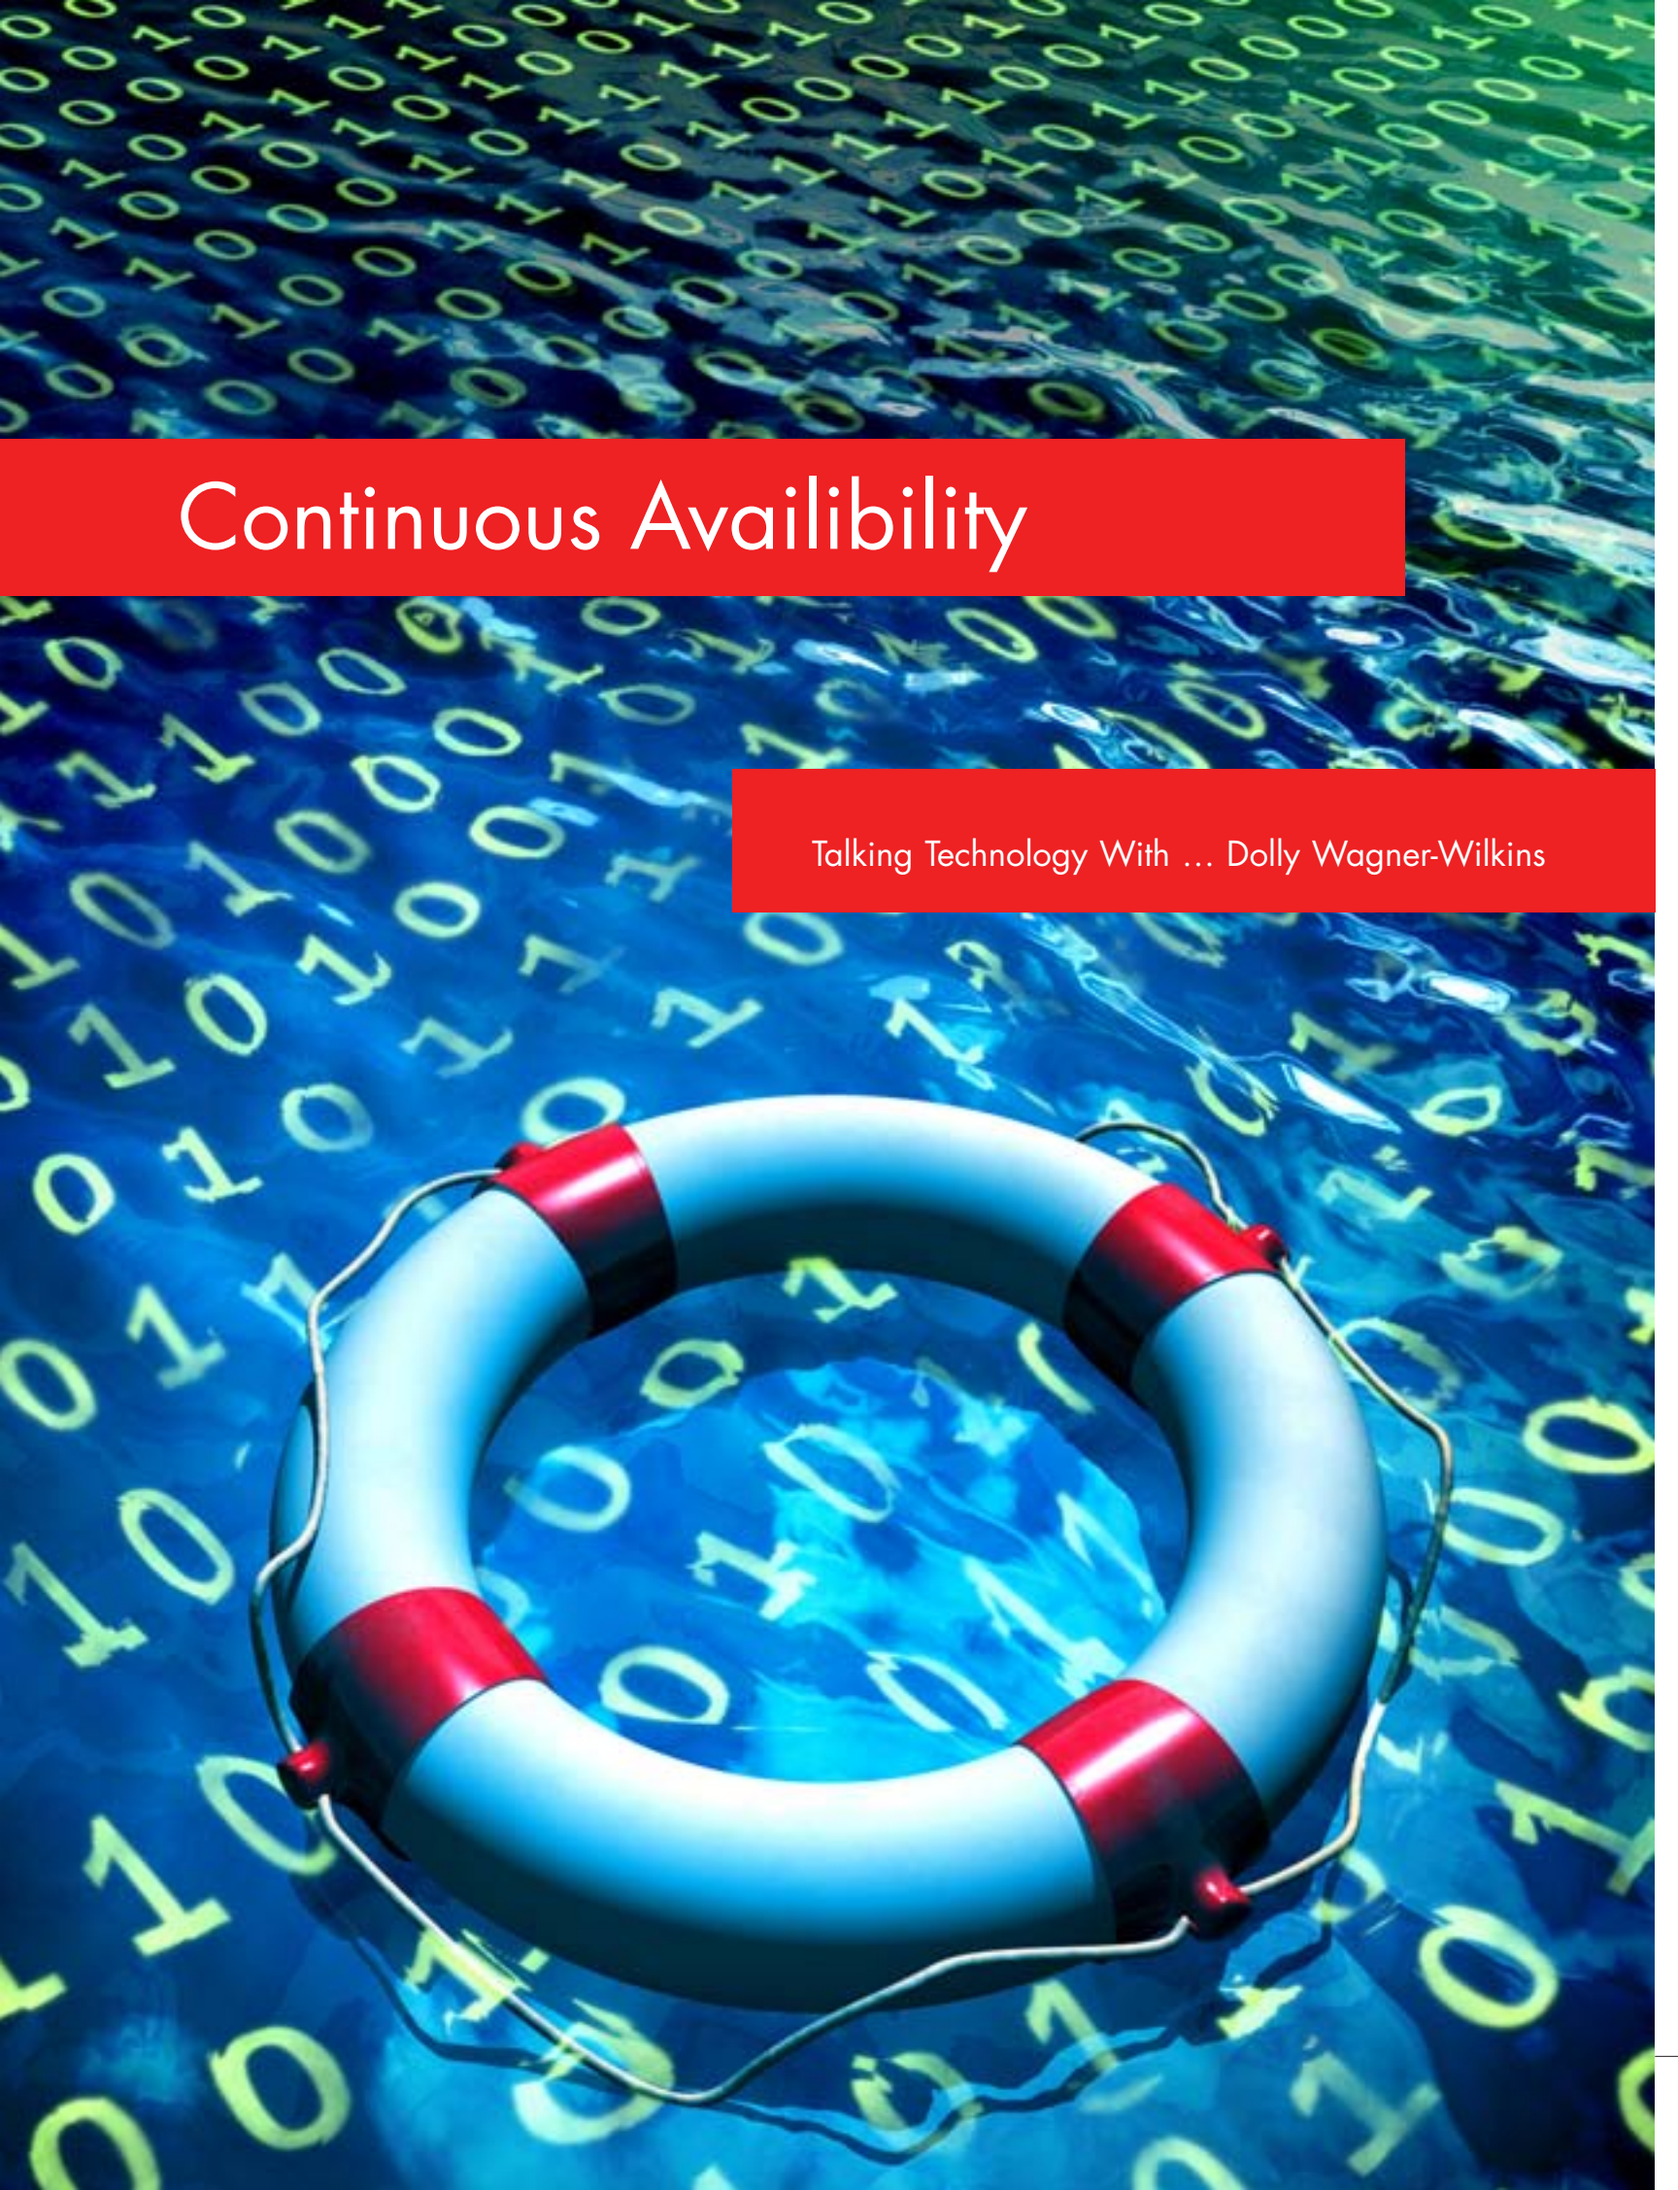

In years past, disaster avoidance was a focal point for *Sabre Holdings*®. But by the early 2000s, that wasn't enough and the company's strategy shifted to disaster recovery. And today, greater needs prevail, calling for a continuous availability solution that better enables disaster recovery.

Do you remember the days when papers were typed on a typewriter, to talk to someone you had to catch them at their desk so they could answer the phone and "Internet" was not yet in the dictionary?

While those days may have been hurried and rushed in some ways, the pace of information and the expectations of customers were entirely different. People understood that tickets had to be printed and printers were finicky, that computers were only so fast and needed to be down for maintenance and that, even in the early days of the "World Wide Web," it was perfectly acceptable to have parts of a website "under construction."

In those days, few companies had even contemplated a disaster recovery strategy beyond data protection, which involved storing a copy of code and data on tapes in an offsite location and hoping it could be recovered if needed. Additionally, airlines were well prepared to operate in "manual mode" with stickers representing seat assignments with a standard supply at airport ticket counters, and reservations agents prepared to write down requests and call customers back to confirm bookings.

As technology evolved in the mid 1980s and early 1990s, our mainframe system was enhanced to run seven days a week, 24 hours a day, no longer requiring downtime for nightly maintenance. Distributed systems were emerging, although these systems were still highly dependent on proprietary operating systems and hardware with limited tools and processes for operating high availability systems. With an increasing reliance on uptime for a wider range of airline operational functions, protecting critical systems against a disaster became a growing concern.

Since mainframe technology was still predominant, the first formal disaster protection plan for *Sabre Holdings* was focused on disaster avoidance, as the idea of running a given system in two locations was still both technically and financially prohibitive.

That said, reliability, innovation and large-scale computing were already expected and part of our corporate culture. Our avoidance strategy centered around a data center that was an underground "bunker," which was flood proof, fire proof, protected from intruders by a retina scan, and able to withstand a Boeing 747 airplane landing on top of it.

While our avoidance strategy was market leading at the time, by the early 2000s, distributed systems were becoming enterprise class, network bandwidth was starting to be measured in gigabytes per second and a petabyte storage array was coming soon. Distributed computing had emerged as a viable mainframe alternative for large-scale computing, and the Internet was maturing as an important distribution channel.

This sparked a whole new pace of business and a new set of consumer expectations. With these advances, it was clear that we had to shift from a disaster avoidance strategy to a disaster recovery strategy.

HIGHlight

> **... it was clear that we had to shift from a disaster avoidance strategy to a disaster recovery strategy.**
>
> *— Dolly Wagner-Wilkins, senior vice president of enterprise operations, Sabre Holdings.*

During the mid 2000s, we implemented our first disaster recovery solution, which was designed to protect airline operations and focused around our core connectivity, reservations and departure control systems. This solution enables a disaster recovery site to become our production site if a significant event disables or destroys our primary data center.

Following the initial implementation, the solution was enhanced to have a full active/active network, linking the primary and secondary data centers with fiber, enabling data to flow between them with unnoticeable latency. This solution is in place today and continues to be at or better than the market. It has a recovery time objective (RTO), the time from the declaration of a disaster until systems are up and running in the disaster-recovery site, of 12 hours and a recovery point objective (RPO), the maximum lag since the last data backup, of 1 hour to 12 hours.

While thankfully it's never been used, the current disaster recovery solution has been successfully tested each year since 2006, proving that in the event of a major disaster, we could operate from our disaster recovery site.

With a solid solution in place, in recent years, we have had the ability to take a step back and think through what a next-generation disaster recovery solution should look like given new and emerging technology and customer expectations.

The dependence on computing technology continues to grow in the airline industry. Electronic ticketing, Web reservations and check-in, and self-service have become the norm, and services that use algorithmic applications to perform rebooking, revenue management and complex shopping functions are available, often doing in minutes what used to be done over hours and days.

As we look at the technical environment today versus the mid 2000s, technology has shifted away from mainframe systems. Most of the systems providing key functionality today are built on our service-oriented architecture (SOA), use robust relational databases that lend themselves to replication, are highly scalable, can be virtualized where appropriate and run in our private cloud.

Network technology is robust with significant bandwidth and features available to minimize latency, add intelligence to load balancing, and run both virtualized servers and networks. Given the flexibility current and emerging technology affords, we believe the next step is really not about disaster recovery per se, but rather about building a continuous availability environment that enables disaster recovery should the need arise.

For many years, we have been running a high-availability environment in a single data center with fully redundant systems. This has enabled us to provide high uptime for our customers, with disaster recovery protection from a major event.

While this design is still the most common in the marketplace, it is vulnerable to localized issues that are significant but far from a total disaster, such as the loss of a portion of the network, database corruption or a significant human error that takes out a single system. With redundancy and good processes in place, these system interruptions are recoverable, but drive longer outages and business disruption than is desirable.

A continuous availability solution extends the high-availability model further by running two high-availability environments for each system across two data centers, closing the gap on restoration time for those larger, localized issues.

This shift in strategy is predicated on several big decisions we've made as we looked to the future of our business.

First, reliability and high availability have been a part of our culture and brand for decades, and we decided to make a significant investment to lead in the marketplace and make continuous availability a priority.

Second, weighing the pros and cons of data center distances, we made a strategic decision to stay in our current data centers that are four-and-a-half miles apart and connected by redundant fiber rings. To enable continuous availability with high-volume systems performing as well or better than they do today, short-run fiber connectivity is a must. When we evaluated the real risk around distance, there was little compelling data supporting further separation.

The majority of incidents that drive the need for continuous availability and even disaster recovery are localized:

- Loss of connectivity or power,
- A hardware failure,
- A bad implementation,
- Human error,
- Flood or fire in a data center.

Some natural disasters can drive a broader impact, but given the low risk of natural disasters in our geographical location and protection against a loss of power with each data center powered by two of three different power grids and full internal power redundancy, we made the decision that closer was better. The benefit of the fiber connectivity and ability to move to continuous availability holds far more value than any real risk of the data centers being located closer together.

Finally, we decided that we needed a much broader and more flexible solution that goes beyond high availability. Requirements included bi-directional disaster recovery so we could recover from a disaster in either data center, giving the ability to fail over a single system, a group of systems or all systems as needed for both true disasters and significant events.

The same solution and investment must enable recovery from a single issue on a single system to the loss of an entire data center, and it must utilize all equipment instead of having idle equipment in a disaster recovery site waiting to be used.

Additionally, we wanted the new solution to be continuously tested, which eliminates the need for an annual disaster recovery drill. In disaster recovery terms, this shifts our midrange RTO from hours to minutes and our RPO to near zero.

This strategy is not without challenges, and since a careful design is crucial, we have applied our most senior engineering and operations teams to the effort. Design elements such as the method of data replication, load balancing across data centers, fail-over methodology and capacity management can easily become complex, operationally challenging and overly expensive if they are not standardized across applications.

Also, discipline is required to carefully apply the wealth of available technology only where it has clear benefit. In short, we believe it's important to strike a balance between capability and simplicity. It's also critical not to assume that technology itself equals a continuous availability solution. Just because a system runs in the cloud, data is replicated across data centers or there is high bandwidth network connectivity does not mean a useable continuous availability or disaster recovery solution is in place.

HIGHlight

> As we move forward with this new mode of operation, we plan to start with a couple of systems, learn from our experience and build our best practices as we grow into the world of continuous availability.

A full solution requires:

- The appropriate technology as well as discipline to keep the technology in sync across the two data centers,
- A testing methodology where the data centers periodically switch between which one is primary for each system,
- A manual or automated fail-over capability at both data center and application levels,
- More advanced capacity planning with careful assessments on capacity requirements,
- Agreements with customers to have connectivity to both data centers.

We will be bringing our first midrange systems live using the continuous availability design by the end of the year, and our plan is to move all of our existing critical midrange systems to this model over the next several years.

As we move forward with this new mode of operation, we plan to start with a couple of systems, learn from our experience and build our best practices as we grow into the world of continuous availability. As our toolbox of design elements and operational practices grows and matures, we will accelerate our pace in moving additional systems to this model. In the meantime, our existing disaster-recovery solution will remain in place.

We will learn a great deal as we work with our technology partners to execute our new next-generation strategy, explore new possibilities and find the best ways to operate in the new continuous availability model.

We welcome your active participation and feedback on our strategy as well as the opportunity to collaborate and share ideas and experiences as we work to maintain and improve upon our best-in-class availability and reliability.

Dolly Wagner-Wilkins is senior vice president of enterprise operations for Sabre Holdings. *She can be contacted at dolly.wagner-wilkins@sabre.com.*