connect
Your Independent HPE Business Technology Community

The Connection

A Journal for the Hewlett Packard Enterprise Business Technology Community

Tuning NonStop SQL Queries to Improve Application Performance

The Dawn of Fault-Tolerant Computing

Building an Integrated Data Management Strategy

May – June 2016 { Volume 37, No. 3 }

XYGATE® Data Protection

All the benefits of HPE SecureData™

No application changes

Simplicity

Secure Stateless Tokenization

Format Preserving Encryption

Page Integrated Encryption

Standards Based AES

Stateless Key Management

Now available from

Hewlett Packard Enterprise

Protect your Data
at rest and in transit

Learn more at
xypro.com/XDP

XYPRO®
Mission Critical Security

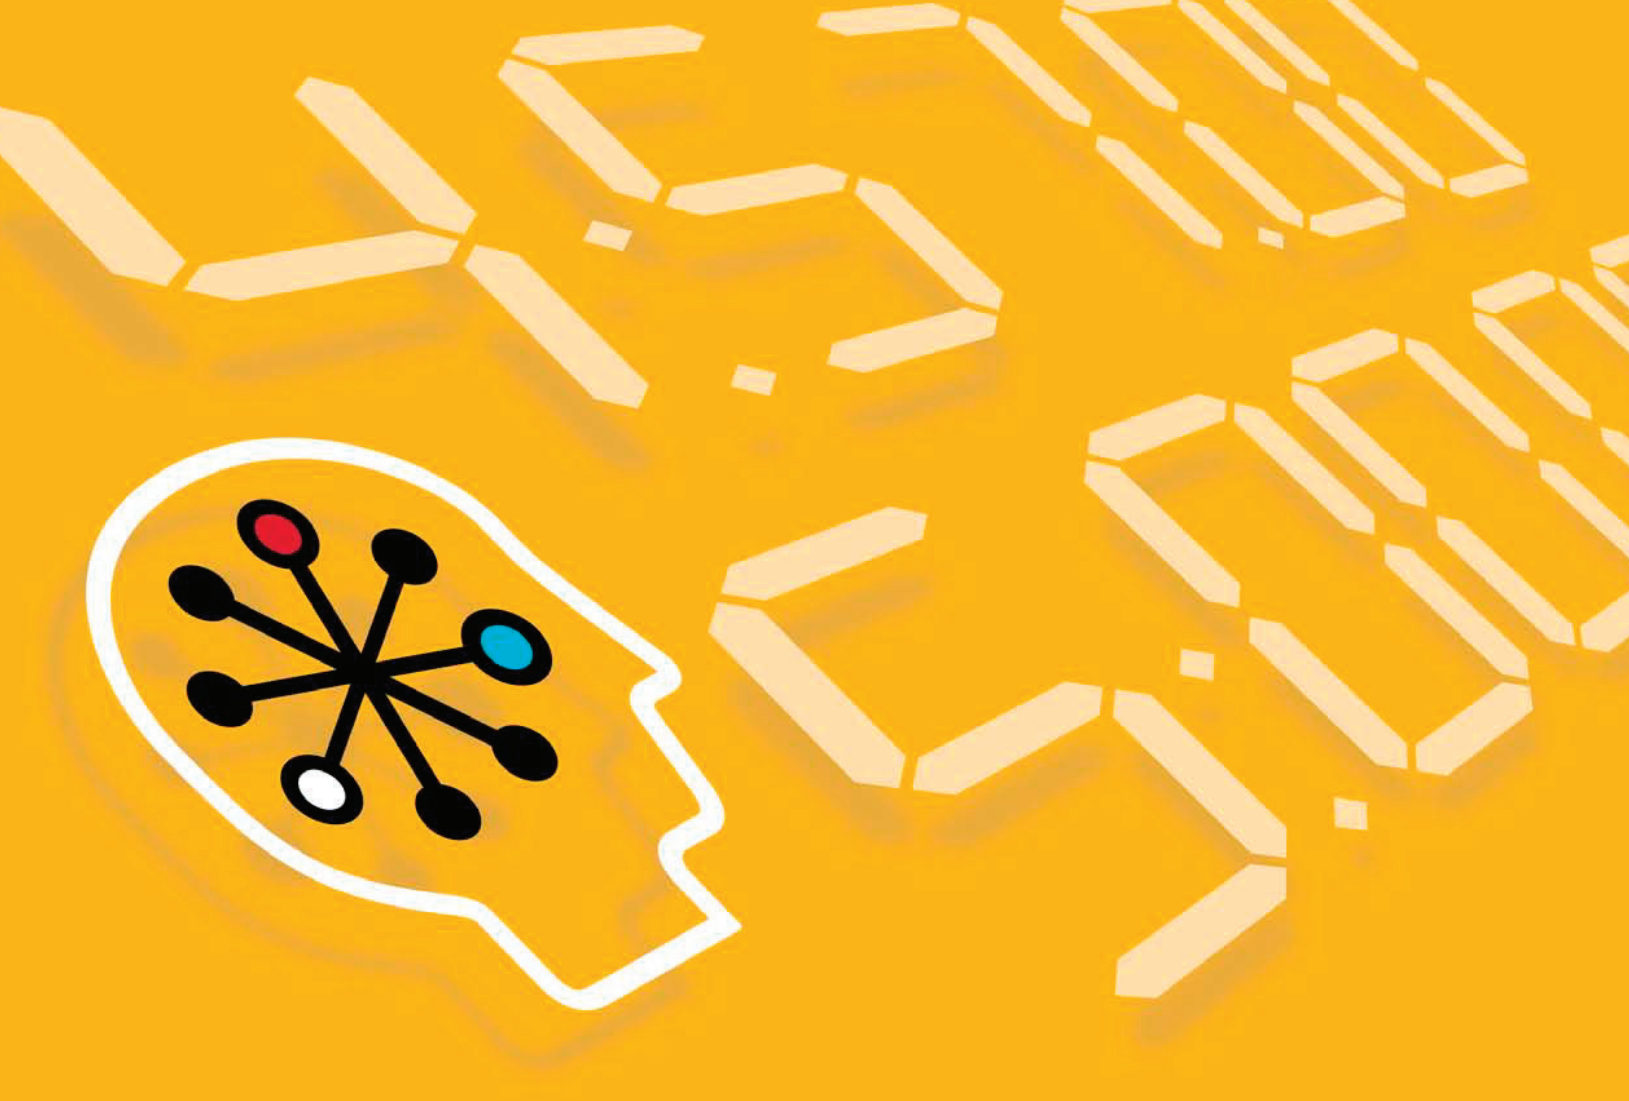

ALTERNATIVE THINKING ABOUT KEEPING TIME:

Guaranteed Time Synchronization

It's 5:00 o'clock at your workstation, but the server next door thinks it's 4:57. How accurate are your transactions?

In today's world of distributed processing and stringent audit requirements, all the systems involved in a distributed transaction must operate at the highest level of accuracy. Without synchronized computer clocks, it's impossible to correlate data from multiple system sources, much less understand transaction flow, pinpoint bottlenecks, or resolve problems.

HP NonStop Time Synchronization provides up to microsecond accuracy across all of your computers, whether you're running HP NonStop, Linux, or Windows operating environments. And, this product is fully backed and serviced by HP.

Technology for better business outcomes.

Key features and benefits

- Dramatically reduces the risk of system events or database records appearing out of sequence.
- Works across your heterogeneous environment — HP NonStop, Linux, and Windows operating systems.
- Automatically synchronizes clocks in a cluster or network; no user intervention required.
- Assures time updates occur safely on all platforms.

Contact your HP representative or partner for a FREE 60-day trial.
Visit www.hp.com/go/timesync


Hewlett Packard
Enterprise
Discover 2016
Accelerating next
Las Vegas June 7–9
Looking forward
to seeing you

Table of Contents

Columns...

OmniPayments

Financial Transaction Switch

The OmniPayments 50%-Less Guarantee

Unlike other switch providers, who base your costs on transaction volume, OmniPayments sells you a one-time, perpetual software license. That is what makes the "50%-Less" difference. No volume fees. No transaction fees. No surprises. No penalties for company growth. All of our prices are presented to you up front. It's just that simple. In addition, our software leasing option allows you to acquire OmniPayments with a minimum initial investment.

- **Comprehensive payments solution for banks and retailers**
- **Built on NonStop for highest possible availability, scalability, reliability, and performance**
- **Proven in production with 14,000 ATMs and 700 million transactions per month**
- **A single OmniPayments system supports up to 10,000 transactions per second (TPS)**
- **Modern, component-based design (BLMs). Open SOA environment**
- **Migrate to OmniPayments. Efficient, pain-free transition. No disruption to customer services**

Affordable OmniCloudX

- NonStop X hosts numerous instances of OmniPayments
- Allows mid-size OmniPayments customers to operate their own high-capacity transaction switches
- Offered on a pay-for-use basis, starts at $5,000/month
- Continuously available with automatic failover to other NonStop X systems
- Complete security functions for encryption-at-rest and encryption-in-flight. PCI-compliant

OmniPayments Fraud Blocker

- Modern and easy to manage
- Preauthorizes millions of transactions in real time and far more effective than complex, compute-intensive competitors
- Sold as part of OmniPayments or as a seamless interface to other providers' solutions

Contact us for a free proof of concept.

OmniPayments Inc.
1566 La Pradera Drive
Campbell, CA 95008 USA
www.omnipayments.com
sales@omnipayments.com
+1 408 364 9915

A Note from Connect Leadership

Wow, the HPE NonStop world is ever growing! I have always been a huge fan of the NonStop community. We are a tight knit bunch. The past few weeks have simply reminded me of just how expansive, yet close, we all are. Over the past five weeks or so I have literally traveled the globe meeting with HPE NonStop customers, many of which I have known personally for much of my career, some I have never physically met until this week. It never fails to amaze me how everyone opens up once we start talking about literally anything we have in common on the NonStop. We all seem to have the same issues regardless of location or application.

This issue is about operations and performance management. This reminds me of my favorite aspect of the NonStop and why I will never respect another platform the way I do this one: Manageability. We do a lot, with less.

Who on the NonStop wears one hat? Of all the people I meet over the years, I find that I can always discuss systems management, database/data management, development, operations, security, performance, batch operations, or any other topic with anyone at any time! I started on this platform as an operator because no one else had the background/experience to do the job. My next job was as a systems programmer for the same reason. Next I was a DBA, then a security analyst. I have always been involved in system performance and security as I assume we all have out of necessity. Eventually we all end up doing every job while someone is unavailable for some reason. What other platform or system can say the same?

What it all boils down to is that NonStop professionals are truly that: Professional. The HPE NonStop server has always attracted the curious and capable and it allows us to be far more effective with far less staff.

So step into your systems management role and enjoy this month's issue. And if you see me around, please ask me how you can help!

Thanks.
Rob Lesan

Rob Lesan

XYPRO

2016 Connect Board of Directors

PRESIDENT
Rob Lesan
XYPRO

VICE PRESIDENT
Michael Scroggins
Washington St. Community College

PAST PRESIDENT
Henk Pomper
Plusine ICT

DIRECTOR
Jamil Rashdi
wasl Asset Management Group

DIRECTOR
Trevor Jackson
SOCAN

CHIEF EXECUTIVE OFFICER
Kristi Elizondo
Connect Worldwide

HPE LIAISON
Janice Zdankus
Enterprise Group, HPE

The Connection

The Connection is the official magazine of Connect, an independent, not-for-profit, user-run organization.

Kristi Elizondo ..... CEO
Stacie Neall ..... Managing Editor
Kelly Luna ..... Event Marketing Mgr.
Keith McLemore ..... Membership Director
Janice Reeder-Highleyman ..... Editor at Large
Dr. Bill Highleyman ..... Technical Review Board
Karen Copeland
Thomas Burg
Bill Honaker
Justin Simonds

We welcome article submissions to the *The Connection*. We encourage writers of technical and management information articles to submit their work. To submit an article and to obtain a list of editorial guidelines email or write:

The Connection
E-mail: sneall@connect-community.org
Connect
P.O. Box 204086
Austin, TX 78720-4086 USA
Telephone: +1.800.807.7560
Fax: 1.512.592.7602

We accept advertisements in *The Connection*.
For rate and size information contact:
E-mail: info@connect-community.org

To obtain Connect membership and *The Connection* subscription information, contact:

Connect Worldwide, Inc.
P.O. Box 204086
Austin, TX 78720-4086 USA
Telephone: +1.800.807.7560
Fax: +1.512.592.7602
E-mail: info@connect-community.org

Only Connect members are free to quote from *The Connection* with proper attribution. *The Connection* is not to be copied, in whole or in part, without prior written consent of the managing editor. For a fee, you can obtain additional copies of *The Connection* or parts thereof by contacting Connect Headquarters at the above address.

The Connection often runs paid advertisements and articles expressing user views of products. Articles and advertisements should not be construed as product endorsements.

The Connection (ISSN 15362221) is published bimonthly by Connect. Periodicals postage paid at Austin, TX. POSTMASTER: Send address changes to *The Connection*, Connect Worldwide, Inc., P.O. Box 204086, Austin, TX 78720-4086 USA.

In the world of business continuity, there's no fire department to call before things get out of control. By then, it's too late. To protect your IT services from fire, or any one of a dozen other serious threats, you need the protections in place *before* the worst happens.

Shadowbase business continuity solutions ensure that no matter how toasty or damp your critical data becomes, there will always be an up-to-date copy available in another location to keep your business online. Don't wait for the fire department to arrive. Instead, contact Gravic today for more information on how Shadowbase software can protect your business.

For more information, please see the Gravic white paper:

Choosing a Business Continuity Solution to Match Your Business Availability Requirements

GRAVIC® Shadowbase

ShadowbaseSoftware.com

Business Partner

Hewlett Packard Enterprise

News from HPE's NonStop Enterprise Division

The NonStop Momentum Continues

2016 is proving to be an exciting year so far for HPE NonStop. The announcement about our strategy to bring NonStop to virtual environments is resonating well with customers and partners. At the same time, sales of NonStop X are increasing with more customers worldwide placing orders and making plans to move to the new platform.

The NonStop community has a rich tradition of coming together to discuss face to face what we are working on and what customers need. This started of course with the old ITUG (International Tandem Users Group) when we were still Tandem Computers, but now has morphed into some 30+ events each year that occur across the United States and around the world. These forums, variously known as Boot Camps, TUGs (Tandem User Groups) or NUGs (NonStop User Groups), provide regular venues for the NonStop community to come together and hear about the latest products and solutions offered by HPE and our partners.

This year so far regional events have had record turnouts with interesting speakers and some great parties. Our own Wendy Bartlett was on hand to talk to customers at MRTUG in Chicago and OTUG in Ohio in March, we had a record turnout of 60 people at our Middle East events in Dubai, Qatar and Saudia Arabia in early April. These are new events in the region, so we were very pleased with the recent attendance. More than 200 people came to hear Andrew Bergholz, Director of Engineering for NonStop take the stage at GTUG in Berlin to talk to European customers about all that's available today and where the future lies. He also gave a well-attended talk that walked people through the architecture and current vision for virtual NonStop. GTUG this year was held at the wonderful Radisson Blu Hotel, which features a very large aquarium over the café and bar so as people met in the lobby to discuss business over drinks they could look up to see very large fish and sometimes a scuba diver floating over their heads.

Besides TUGs and NUGs, NonStop was at Mobile World Congress this February with a live demo of our lab prototype for virtual NonStop. We demonstrated NonStop running in KVM virtual machines on Linux servers with Telco applications. Virtual HSS, INS and UDR software from the HPE CMS organization were shown on the system as well as Virtual Local Number and Virtual Data Service products from Inovar Software. We were one of the few booths at the show with a live demo. Mark Pollans, a senior Product Manager, had the pleasure of presenting the demo to Meg Whitman who dropped by to see it after others recommended it to her.

But of course that's not all. In May there's the N2TUG meeting in Dallas, where word on the street is that Jimmy Treybig may be on hand to check in with the NonStop community. At the end of May we are hosting our first ever Symposium for the NonStop Partner community here in Palo Alto, California. We're very excited that so many companies are planning to attend and looking forward to seeing everyone. The week following of course is the HPE DISCOVER 2016 show in Las Vegas, where Randy Meyer, Jeff Kyle and Mark Pollans will be on hand to welcome you and share further demos of NonStop's latest capabilities. Following DISCOVER 2016 there are two events planned in Asia, KNUG (Korean NonStop Users Group) on June 21st in Seoul and Taiwan NonStop Technology Day on June 23rd in Taipei. Other regional events being planned include VNUG (Viking NonStop User's Group) held in Scandinavia, OZTUG being planned in Australia, and CTUG (Canada Tandem User's Group) planned for early October. Before you know it, it'll be November and we all be together again at the 2016 NonStop Technical Boot Camp in San Jose November 13th through the 16th.

I have had Product Managers from other teams ask me how we do it, how we keep so many events going each year around the NonStop platform. I tell them that it starts with having interesting new topics, but these events are a testament to the Connect team and their liaisons, working with our NonStop Partner and Customer community and the NonStop group at HPE. Together we are what keeps NonStop events well attended and flourishing.

As this edition of Connection Magazine goes to press, we'll be only a week or two away from announcing a Carrier Grade version of NonStop X and completing upgrades on both NonStop i and NonStop X to RoHS compliant hardware. This helps round out the NonStop X product line and solidifies the future of both product lines. On the software side, look for announcements about Java 8 and of course about YUMA which is soon to be an available as a product – under a new, official name, NonStop Application Direct Interface.

Meanwhile, we encourage you to cozy up with this issue of Connection Magazine, by the pool with a cold beer or frosty iced tea and enjoy articles aimed at the Database enthusiast such as Part 2 of Frans Jongma's article on Native Tables and SQL/MX and John Furlong's article on Tuning NonStop SQL Queries to improve your application's performance. Justin Simonds is offering an interesting view of where the Healthcare industry is going with his article Toward Precision Healthcare and Bill Highleyman is providing a fascinating view of how Fault Tolerant computing began, the different approaches that were tried and how we got to where we are today. The issue is rounded out with an article from Glenn Garrahan on using a Policy Based Data Management approach to the data you need to store and Richard Buckle's entertaining views on how to think about Managing NonStop systems.

We hope you enjoy this edition of Connection Magazine and we will be happy to see you at the Discover show in Las Vegas or one of the many other NonStop events this summer and fall!

Karen

Karen Copeland
Manager, WW NonStop Product Management
Mission Critical Solutions
Hewlett Packard Enterprise

On The Road to MENUG and GTUG

Kristi Elizondo >> **Chief Executive Officer** >> Connect Worldwide

This year Connect Worldwide and Hewlett Packard Enterprise presented three full days of NonStop product updates, roadmaps and networking on April 10, 2016. The MENUG Chapter meetings were hosted in three different cities. Our first stop was Riyadh, Saudi Arabia. The event was at the Intercontinental Hotel overlooking a nine-hole par 3 golf course that is fully lit at night.

The road show focused on Security and Business Continuity. All the major Saudi Arabia banks were represented. The design of the meeting was different than previous chapter meetings. The meeting started with the HPE NonStop road map and updates, followed by short vendor presentations. After lunch, the vendors conducted one-on-one's with customers. This process, similar to speed dating, allowed for all of the customers to meet individually with the vendors. It was successfully received by attendees.

Our next stop was Doha, Qatar. The meeting was held at the Doha Grand Hyatt on the Persian Gulf. The last stop in the Middle East was Dubai, UAE. We met at the HPE office in Internet City. We chose to Uber to the office and it was the first time any of us saw Uber Chopper as a menu option. You know you're in Dubai when you are offered a helicopter ride to work! The day was overcast with sand storms so we opted for the traditional ride. It is important to note that it rains five days out of the year in this area of the world. It rained in Doha and Dubai, so we think they may welcome us back! A special thanks to Grant Adonis, Clive Abrahams, Esther Sanchez, and Ad Kiljn for participating in the core event team. Thank you to our sponsors, Comforte, ETI-NET, Gravic, Idelji, XYPRO, and Hewlett Packard Enterprise.

While on that side of the world my final stop was Berlin, Germany. Connect Germany and GTUG held their annual IT-Symposium on April 18-20. The conference was well attended with over 200 delegates. They ran 5 tracks in both German and English covering mission critical systems. Hosted by the partner community, and staying true to German tradition, the chapter party cruise on the river Spree was a night to remember.

In 2017, the European event will be hosted by BITUG in London in the spring. Stay tuned for further details.

Thanks for reading, and if you are interested in attending any of the chapter meetings please find locations and dates on the Connect calendar. www.connect-community.org. **Come join us!**

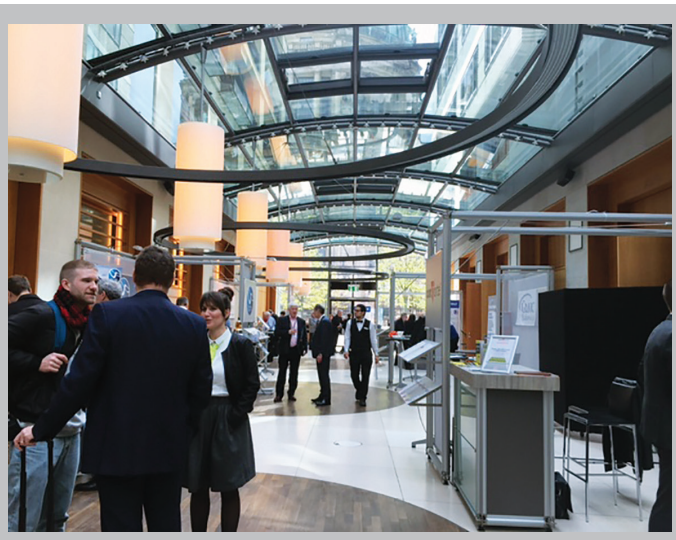


XYPRO
XYPRO

Let Your Voice Be Heard Via the New Connect Website

Dr. Bill Highleyman >> **Managing Editor** >> Availability Digest

The new Connect website is now online, and Advocacy issues can be submitted directly from the home page. This is a marked improvement over the earlier website, in which the Advocacy submission form was well-hidden. Let's go through the steps of using the Advocacy services on the new Connect website.

First, Log On

The first step, of course, is to log on to the website. Warning – Your old credentials will no longer work. If you are already a Connect member and have registered on the new site, you have your user name and password. If you are not a Connect member or have not reregistered on the new site, registration is simple. Plus, it is free for HPE customers and employees. To access the registration page, click on the ">" symbol on the right of the top graphic until you get to it. Then click "Become a Member." You will be asked for some personal and professional information. Use your email address as your user name, and select a password.

Send Us Your Advocacy Submission Form

Once logged on, click on the "Community" tab at the top. You will get a list of topics, one of which is "Advocacy – Community Voice." Another is "Advocacy Submission Form." Click on this link.

All you need to do is to enter your name, your email address, and your issue. You also may enter your company, your title, and your phone number. There is a drop-down list with a series of subjects. Use the list to further categorize your issue. Subjects include NonStop, OpenVMS, Superdome, Oracle, Storage, Networking, Security, Software, and Desktops/Laptops.

Clicking the Submit button will forward your issue to the Connect Advocacy Committee for review.

Connect's advocacy program is intended to help members submit their product or service enhancement requests through the Advocacy Committee. Each request will be logged and reviewed by the Advocacy Committee, which meets monthly. Your request will be passed to HPE for comment and/or action, and HPE's response will be returned to you. The response may be that they will consider the request, will act upon the request, or will take no action on the request.

The Advocacy program is not intended to replace HPE's support mechanism but rather to provide Connect members a collective voice on product enhancement concerns.

The Connect Advocacy Program

The "Advocacy – Community Voice" link takes you to a description of the Connect Advocacy program. As the description points out, the purpose of Connect's Advocacy program is to leverage the voice of the Connect global community to deliver unified messages to Hewlett Packard Enterprise. The Advocacy program's ultimate goal is to facilitate a two-way stream of communication that results in solutions and positive changes for the benefit of both HPE and the user community.

Other Ways to Submit Issues

Issues can also be submitted through Connect SIGs (Special Interest Groups) and Connect chapters.

SIGs

Connect SIGs provide opportunities for members to share best practices, learn from subject matter experts, and connect on areas of special interests via online discussions, webcasts, Connection articles, and face-to-face sessions at Connect events. The SIGs also serve as a means to advocate the needs and interests of members' companies to HPE or HPE partners. Each SIG is headed by a leader who is responsible for ensuring communication between the SIG members and HPE.

Connect currently is supporting the following SIGs:

- Enterprise Networking
- Enterprise Storage
- Enterprise Security
- Cloud
- Converged Systems
- NonStop
- OpenVMS
- UNIX/HP-UX
- Linux

Advocacy issues submitted during a SIG meeting or via other forms of communication will be passed to the Advocacy Committee for referral to HPE.

Chapters

Connect chapters provide members with the opportunity to meet face-to-face, network, and discuss industry topics of interest in a local venue. Connect currently has thirty-six chapters in North and South America, EMEA, and the Asia/Pacific.

Advocacy issues submitted during a chapter meeting will be passed to the Advocacy Committee for referral to HPE.

Let Your Voice Be Heard

The Connect Advocacy program provides a variety of ways for you to discuss your issues and concerns relating to HPE products and services with other users and with HPE. Please take advantage of these avenues to improve the HPE experience for us all.

Dr. Bill Highleyman is the Managing Editor of The Availability Digest (www.availabilitydigest.com), a monthly, online publication and a resource of information on high- and continuous availability topics. His years of experience in the design and implementation of mission-critical systems have made him a popular seminar speaker and a sought-after technical writer. Dr. Highleyman is a past chairman of ITUG, the former HP NonStop Users' Group, the holder of numerous U.S. patents, the author of Performance Analysis of Transaction Processing Systems, and the co-author of the three-volume series, Breaking the Availability Barrier.

NonStop Innovations Deep Dive

Striim and the Future of Streaming Data and Analytics

Gabrielle Guerrera >> NuWave Technologies

As tends to be the case in technology, things change quickly, and Striim (formerly WebAction) is no exception to the fast-paced nature of the industry. After undergoing a recent rebranding, the company is in full swing promoting their unique platform for real-time data streaming and analytics.

I recently had the opportunity to chat with Sami Akbay, Executive Vice President, about all the changes the company has undergone and what it has in store for the future.

Gabrielle: **Hi Sami, I've interviewed you and some other members of your executive team for The Connection before, but the company was called WebAction back then. Could you recap the company background for those who aren't familiar?**

Sami: WebAction, as we were formerly called, was founded in 2012. We specialize in streaming data integration and analytics in real time, which helps the user make sense of large amounts of data as it is created, and also helps them organize it in ways that are useful to their company.

Essentially, what our product does is it processes data from many sources and makes sense of it in real time, then it retains only the information that is interesting to the end-user, as defined by the user. The volume of data that most companies have access to is very large, but there is only a small amount that is useful. Our product finds that useful information and puts it back onto the NonStop.

Prior to the founding of WebAction, we were certainly not new to the NonStop space. Our company was founded by some of the core leaders of GoldenGate Software (which was acquired by Oracle in 2009), including our CEO, Ali Kutay. He was also the chairman and CEO of GoldenGate, and I used to run product management and marketing for GoldenGate before it was acquired by Oracle. So we are definitely not new to the software industry, nor to the NonStop, and we are happy to be able to bring our experience to help other companies get the most out of not only their NonStop server's capabilities, but to also have a better understanding of their data.

Gabrielle: **What prompted the name change from WebAction to Striim?**

Sami: We got a lot of feedback from people who said that the name of the company didn't do a very good job of describing what we had to offer, which is a really unique platform. We took some time, did market research, and decided that the name Striim (pronounced "stream") would be an interesting way to describe what we did.

In our name, the first "i" stands for "integration". Our platform takes large amounts of data from multiple sources and can deliver them to many types of systems extremely quickly. The second "i" stands for "intelligence". As the data comes in, we are able to provide real-time analytics about the data so the user can easily understand the information that is being created. Striim uniquely combines both streaming integration and streaming intelligence into a single platform. It helps put large amounts of high-speed data in context, which is really important for understanding customers or industry changes.

Gabrielle: **Has there been any confusion since the name has changed?**

Sami: No, not really. We've received a lot of positive feedback from people, and we reached out to many of our friends in the industry to get feedback before we changed our name to see how

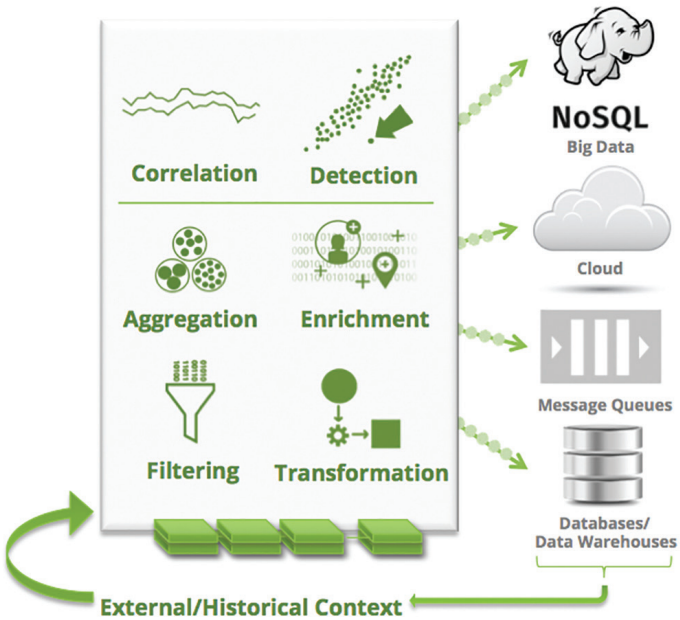

people would feel about it. There hasn't been a lot of confusion because what we offer is very unique, and the name Striim does a much better job of explaining that than WebAction did. Aside from someone occasionally mispronouncing it, the rebranding has been extremely positive.

Gabrielle: **Did you change any of your product offerings when you rebranded?**

Sami: Our core Striim Platform focuses on improving how companies collect and use their streaming data, and that has stayed the same since we changed the name. Striim helps companies leverage a wide variety of different types of streaming and operational data the instant it's generated, before it lands on disk. From transaction or change data from Oracle, MySQL and HPE NonStop, to log files, social streams, sensor data, and much more. Then, this data can go through a context loop (see diagram on previous page), where results can be further analyzed to uncover additional business insights or fed back into the data pipeline to give context to streaming data.

As I mentioned before, many of the companies who would use our product, and especially NonStop users, are dealing with extremely large amounts of data. However, only some of that data is really useful. Our platform is able to filter, aggregate, transform, and enrich data in-motion so companies can see what is happening as it happens. The user can correlate data across multiple streams and set alerts to detect outliers or trigger specific workflows. And all of this is easily managed and viewed through our easy-to-use dashboards. They can also load their streaming data to any target, such as Hadoop, Kafka, AWS, NoSQL databases, or their own data warehouse. We make it as easy as possible for our customers to use our product, and that hasn't changed at all.

Really all we did was make it more clear to people who we are and what we do, and we have been really pleased with the feedback and the reaction that we have received from people since we made the change.

Gabrielle: **Do you have any new products that you're working on?**

Sami: The nature of software is that it must change constantly to keep up with market demands and new technology. We are constantly making updates to improve what we offer to ensure that our customers are able to take full advantage of their data, as well as keeping up with new technology and trends. So while we are not releasing new products frequently, we are on a cycle of constant updating and offering new, better releases of our core platform. Our most recent versions are always available to our customers. However, we are always staying on top of the latest technology, so we constantly research new product options for the future.

Gabrielle: **Do you think that technology like this will have an impact on the future of NonStop?**

Sami: NonStop is all about real time, and so are we. I think that as technology continues to progress, having access to your data and being able to process and understand all of it immediately will only continue to get more and more important. If you have no way of analyzing your data quickly, it can get buried, and digging out that buried data can be time consuming and costly, especially when you are dealing with mission-critical information, as most NonStop users are. This is particularly important in the financial space, where many of our NonStop customers are operating.

Everyone is lining up to get LightWave Server

Don't get left behind.

LightWave Server™ uses JSON and REST technology to send data to modern clients like mobile apps and browser-based applications that run on virtually any platform.

Learn more at
www.nuwavetech.com/lwsconnection

NuWave TECHNOLOGIES

We are also very partner-friendly, so we have a lot of partner companies that integrate our platform into what they have to offer to ensure that their customers are getting the best streaming data processing and analytics possible.

Gabrielle: **Who are some of your partners that you work with?**

Sami: We partner with a variety of different technology companies and services providers to bring data integration and streaming analytics to a variety of companies in many different industries. We work with companies like Confluent, Ericsson, Huawei, Hortonworks, MAPR, Amazon Web Services, Cloudera, Southpartner, BIAS, GTEK, Microsoft Azure, Snuvik and many others. We are always open to new partnerships with companies who have customers in need of exceptional data streaming.

Gabrielle: **What are the main industries you find yourself focusing on, both inside and outside of the NonStop space?**

Sami: As I mentioned, we work a lot in the financial space, which is where we find a lot of our NonStop users. Our other big focus is telecommunications. In both industries, there is high data velocity and volume and the time-value of the data (see graph) is very high.

Gabrielle: **Your company has undergone a lot of exciting changes in a very short amount of time. Where do you see Striim going in the future?**

Sami: Streaming data and the ability to integrate and analyze real-time information is very important. It will only continue to grow and become even more of an asset to companies in the future, so we are really excited about what the future holds for Striim, as well as our customers as they continue to grow.

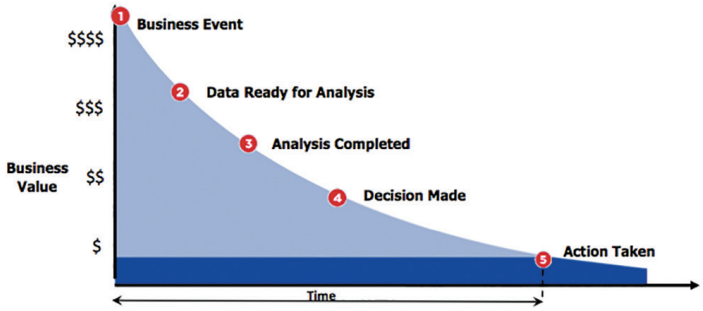

We do have some really exciting things in the works for the coming months, especially with some new customers in the NonStop space that we are looking forward to working with.

Our goal is to provide our customers with the ability to understand and organize their data as quickly and easily as possible, and we only see that continuing to become a strong need with more capabilities in the future.

Gabrielle: **Thank you, Sami. I look forward to following Striim and seeing what other surprises you have in store.**

Sami: It was my pleasure, thank you. We are very excited for where the company is going now as Striim, with an even clearer focus and mission, and for what we can provide for our customers and their data.

Check out the follow-up article on the NonStop Innovations blog at www.nuwavetech.com/hp-nonstop-innovations.

Gabrielle Guerrera is the director of business development at NuWave Technologies, an HPE NonStop middleware company. She is also the co-author of the NonStop Innovations blog, which takes an unbiased look at the latest innovations and announcements in the NonStop space. She is an MBA candidate at Babson College in Massachusetts.

Toward Precision Healthcare

Justin Simonds >> **Master Technologist** >> Americas Enterprise Solutions

This past October (2015) the 9th annual Sino-US Symposium on Medicine (http://med.stanford.edu/sino-us.html.html) took place at Stanford University. I was fortunate to receive an invitation to this event. This symposium is a conference between top medical professions and institutions from China and the United States. The purpose is to share information on Healthcare. The theme of this conference was 'Big Data and Healthcare', although "Big Data" was a means to an end, which was getting to the main conference concept – precision healthcare. The conference was much more about medicine and treatments than it was on analytics but it was fascinating even for a non-doctor. There were many excellent speakers including Nobel Laureate Michael Levitt, PhD.

The talks were quite broad in scope - cardiovascular disease, various cancers, diabetes, brain abnormalities, and skin disease along with many other topics (go to the website for the complete agenda). I thought one of the speakers made a profound point when he started his presentation with the observation that until recently what they had all been practicing was sick-care but now, with the advances in genomics, they were finally at the very beginning of practicing healthcare. Many of the speakers were studying genomics and its effect on disease, especially chronic disorders. The conference confirmed what I had been reading that eight out of 10 older Americans are faced with the health challenges of one or more chronic diseases (http://www.cdc.gov/aging/index.html). Chronic diseases are responsible for 60% of deaths worldwide and account for three-quarters of America's direct health expenditures (http://www.who.int/features/factfiles/chp/01_en.html). People with chronic diseases cost 3.5 times as much to serve compared with others, and account for 80% of all hospital bed days and 96% of home care visits. So work in this area benefits the population in general in terms of health and quality of life but also has profound financial impacts as well. As readers know NonStop was involved in healthcare through GE healthcare and the Centricity Enterprise product. With GE exiting that area of the marketplace NonStop's largest healthcare ISV is ironically at end-of-life. Many people, myself included feel healthcare is a vertical in which NonStop should participate. Based on loss of life alone it makes sense for NonStop to be a core platform for mission critical, let alone life critical applications but considering the additional financial implications of this market, NonStop should be required.

The work of many of the doctors presenting at the conference was around the pursuit of genomic mappings for disease. Although some amazing work has been accomplished we are clearly on the threshold of this effort. The original genomic mapping was commissioned in 1990 and was expected to take 15 years. Thanks to computer advances a rough mapping had been completed by 2000 and was declared complete in 2003. So roughly 13 years and a few billion dollars to map the first genome. At the conference they said mapping could be done today in 48 hours for $1,000.00. That is amazing progress and makes the mapping and testing of patients both practical and affordable. And of course the speed will continue to increase as the price continues to fall. This is important and shows that the healthcare industry is embracing technology in very new and interesting ways, most especially concerning analytics. The Healthcare industry has traditionally been a laggard in terms of technology adoption. Big Data and the Internet of Things is having a pronounced effect on technology adoption. We are seeing the healthcare industry start to embrace technology in three main areas; data collection, data sharing and analytics. Data collection began in earnest with the establishment of EMR (Electronic Medical Records). In 2004 HIMSS (Healthcare Information Management Systems Society) published EMRAM (Electronic Medical Records Adoption Model) which was a framework for measuring the industries advancement and use of electronic records. Now that the industry has been collecting these records a method for sharing needed to be established. Obviously information sharing was the theme of the conference I attended but now we are speaking of actual data sharing of patient information. The proposed HIE (Health Information Exchanges) are charted to do specifically that. However, there does not exist anything similar to EMRAM to track advancement and usage. The bigger problem with HIE is funding. Who pays for these exchanges? Since the transfer is primarily medical information not a financial transaction. Getting someone to foot the bill for the infrastructure and transfer remains an issue. This brings us to my favorite, analytics. We've collected information, now how do we use it to make Healthcare a data driven enterprise and convert from a sick-care industry to healthcare? Interestingly an analytic framework for Healthcare has been developed. It really took the framework for EMRAM as the basis and is more a methodology for tracking the advancement and use of analytics. It is a nine level model (see figure 1).

Healthcare Analytics Adoption Model

Figure 1

As we go up the levels of the model more information is added at each level. In addition less and less latency can be tolerated. Level 8, tailoring patient care, if we begin talking about wearable technology and sensor data, that data will need to be processed in near real-time. Like OnStar reporting information on vehicles, most of the information is repetitive and predictable but every now and again special sensors go off – accident alert, air bags deployed, etc. Similarly we hope most information delivered from a remote patient is repetitive and predictable, but not always. There was an old HP video called "Cool Town" and it lives on at https://www.youtube.com/watch?v=U2AkkuIVV-I. Toward the end there is a dramatization of a healthcare emergency and response which, looking at the new wearables was quite prophetic. It does demonstrate possibilities but only if information is constantly being streamed and analyzed in real time.

For those that might have seen my Internet of Things presentation you might remember my car accident use case. A driver trying to avoid a child running into the street swerves and flips his vehicle. He and his two passengers need medical attention. In the not too distant future alerts from sensors are sent out and directed to 'smart city' 911 services. A first responder is alerted and provided with up to the second best routing to get them to the accident. Smart street lights can flash along the proper route and even change color as the first responder vehicle approaches (smart cities). While the first responder is on the way medical information is collected from wearable technology on the three victims and transmitted to the first responder and to the hospital where the driver and passengers will be taken. In the hospital medical history is also acquired, family contacts and primary care physicians are alerted to the situation. Health insurance information is collected and begins pre-registration of the driver and passengers at the hospital. The first responder arrives and while the accident victims are being transported high resolution cameras allow the emergency room staff to get their first look at the injuries and develop a plan and to assemble a medical team before the patients arrive. How much of this technology already exists for all of this today? Most of it. What doesn't exist is the communication infrastructure and standards for sharing all this information. But it's coming. And, in my example, which parts in the above use case are critical?

The patient data collection through wearables or sensors is all about a new term being used in healthcare called Remote Patient Monitoring (RPM). "Remote monitoring is extremely important and probably in the forefront of mobile technologies now," explains David Lee Scher, MD, FACC, a Harrisburg, Pennsylvania cardiologist and mobile-tech consultant, "because of its potential importance in decreasing hospital readmission rates, which are a big headline because they are responsible for penalties that the Centers for Medicare and Medicaid Services (CMS) is now imposing on hospitals that have readmissions within 30 days for certain diagnoses." (see: http://medicaleconomics.modernmedicine.com/node/384265?page=full). At the conference many of the studies that were being started involved remote monitoring and data collection. For example MYHeart is an application developed at Stanford to collect information from the general public. Only available with Apple at the moment but take a look at https://med.stanford.edu/myheartcounts.html.

In the future we can get a sense of the combining of one's genetic information with day to day (second to second) monitoring. The future view is the capability of detecting very early warning signs based on your genetic predispositions for certain diseases. Imagine getting a 5 year warning about your heart with a list of what to do to avert it?

NonStop will continue to have an important role in this exciting area. NonStop OS has always been a perfect match for OLTP type applications. The remote patient monitoring – regular, structured information coming in almost continually is not dissimilar to OLTP. Additionally NonStop has always excelled at message-switching type applications. As the flood of information starts coming in from a patient demographic NonStop can be used to process highlighted/emergency information. Additionally this information can be switched in near real-time to backend systems for additional processing. NonStop supplies a high degree of security in that the vulnerability list is extremely low which is perfect for a very concerned healthcare industry consumed with privacy regulations. As the Internet of Things rolls out in healthcare NonStop continues to be the preferred platform for the Internet of mission-critical things that require immediate uncompromised attention.

As the healthcare industry moves toward digital healthcare and as consumer based/patient sensors become the norm (smart pacemakers, insulin pumps, dialysis, etc.) more and more analytics will be required to understand what is normal and what constitutes a pending emergency. The Hewlett Packard Enterprise Big Data Reference Architecture (BDRA http://h30507.www3.hp.com/t5/Servers-The-Right-Compute/The-HP-Big-Data-Reference-Architecture-It-s-Worth-Taking-a/ba-p/6795401?jumpid=reg_r1002_usen_c-001_title_r0002#.Vi2zIZjruM8) will become important for the deep analysis, genetic profiling and social trending to understand patients. NonStop will be required for the "911" alerts identified by these deep and predictive analytics.

Justin is a Master Technologist for the Americans Enterprise Solutions and Architecture group (ESA), a member of the HPE IT Transformation SWAT team, and a member of the Mainframe Modernization SWAT team. His focus is on real-time, event-driven architectures, business intelligence for major accounts and business development. Most recently he has been involved with modernization efforts, Data Center management and a real-time hub/Data Warehouse system for advanced customer analytics. He is currently involved with Labs on several pilot projects. He is currently working on cloud initiatives and integration architectures for improving the reliability of cloud offerings. He has written articles and whitepapers for internal publication on adaptive enterprise, TCO/ROI, availability, business intelligence, and the Converged Infrastructure. He is a featured speaker at HPE's Technology Forum and at HPE's Executive Briefing Center. Justin joined HPE in 1982 and has been in the IT industry over 34 years.

Tuning NonStop SQL Queries to Improve Application Performance

John Furlong >> Merlon

Introduction

Query tuning is the process by which you answer the question "How can I make this query run faster?"

There are many possible reasons why an SQL query runs slowly.

- The system may simply not have enough processing power, memory, or disk bandwidth to handle the application load.
- There may be a problem with the system configuration, such as insufficient disk cache, MX buffer space, or kernel-managed swap file space.
- The files used to store the database tables may be badly fragmented.

Such "system-level" problems are outside the scope of query tuning, and must be addressed separately before query tuning can be effective.

The query tuning process treats the query text as the sole statement of the problem. Its goal is to produce an execution plan that runs the query faster, while returning the same set of rows as before.

Why queries run slowly

From a business process perspective, most reasonable queries will return a moderate number of rows.

For a reasonable OLTP query, the maximum number of rows returned is around one hundred or so (and is often much less than this). The user of an interactive query doesn't want to browse through many more rows than this.

Even for a reasonable batch or reporting query, the user does not want to browse through more than a few hundred (or very few thousand) or so rows. If a batch-type query returns more than a few thousand rows, it needs to be broken down into multiple more closely filtered queries, perhaps aimed at multiple classes of user.

Modern NonStop systems are quite capable of executing a query that returns a few hundred (even a few thousand) rows in a reasonably short time. So the question arises; Why does a query that returns a modest number of rows take too long to run?

One answer is that the query accesses and processes too many rows that are eventually discarded, and don't contribute to the final query result. The goal in this case is to avoid performing such non-useful work.

Another such class of query is one that returns a reasonable number of rows, but which is required to aggregate a very large number of rows in order to produce the required results. A query containing a GROUP BY clause is a common example of this. These queries can be difficult to tune. A common solution is to pre-aggregate the data in the database, and to query the summary data without accessing the details. The careful use of triggers can be used to synchronize the redundant summary data.

The Query Tuning Process

Query tuning is a multi-step process that generally goes something like this:

1. Establish a baseline. Generate the execution plan and establish the current performance characteristics of the existing query. This gives you a base for comparison when you perform tuning experiments.
2. Analyze the current execution plan. You should understand the path to the data that is being taken by NonStop SQL.
3. Identify problems with the existing plan, and come up with a better plan.
4. Make changes to get the plan you want. This is not always obvious. You may have to perform multiple tuning

experiments to arrive at the solution.

5. Having implemented the changes, generate a new execution plan, and re-measure the performance.
6. Compare the results with the baseline. If things have improved, incorporate the change into the new baseline, otherwise abandon the change.

Repeat steps 2 through 6 until you are satisfied with the query performance.

Understanding Execution Plans

An SQL/MX execution plans sets out the steps that the SQL executor performs when running the query. Each step is represented by an operator which processes the data in some way. The operators are connected together in a tree structure. Typically, data flows from the leaf operators in the tree up towards the root operator, which delivers the results to the application.

An SQL/MP execution plan is presented as a sequence of steps. Each step has one or more operations. This apparent simplicity can be deceiving however. The description of steps and operations can contain details that override this apparently linear sequence. As it turns out, an MP plan also has a tree structure. For example, step 2 might be a right child or a parent of step 1 depending on its "characteristic", whereas a sort operation may be a right child, left child, or parent of its step, depending on its "purpose".

An execution plan contains a wealth of details which may make it seem complicated. However, you can achieve a good understanding of a plan by concentrating on a few essential elements.

- Join Order. This is the order in which the tables in the query are joined together. This is the hardest problem that the SQL compiler must solve. Moving well-filtered tables toward the start of the join order will often help performance.
- Join Method. There are three fundamental join methods, hash join, nested join, and merge join. Each method has its advantages and drawbacks. (Although MX supports 16 different join operators, they are all variations on the three basic types.) Hash joins typically have lower per-row costs, but may access rows that are subsequently discarded. A hash join does not scale well when the hash table overflows to disk. Nested joins usually have higher per-row costs but tend to access only the rows required. They scale better than hash joins. Merge joins work quite well when their inputs are already sorted in the required order, although a hash join will usually out-perform a merge join.
- Access Paths. An access path is the route taken to the data required by the query. Data is accessed by a scan operator. There are a number of scan operator variants including unique, subset, and full, which access base tables or indexes. There is a close relationship between join methods and access paths. For example hash joins tend to use subset of full scans, while nested joins tend to use unique or subset scans. The per-row cost for a unique scan is higher than that of a subset or full scan, but typically only targets required rows.
- Parallelism. MX supports three types of parallel processing in a query: partitioned, pipelined, and independent. MP supports partitioned parallelism. Using parallelism can improve performance for queries that process moderate to large number of rows. However, it can degrade performance on queries that process only a few rows.

The output from an SQL explain command does a good job of describing the details of the plan operators, but doesn't help with the bigger picture. You can't easily visualize the shape of the plan, and how the operators fit together. To do this, you need to view the plan as a tree.

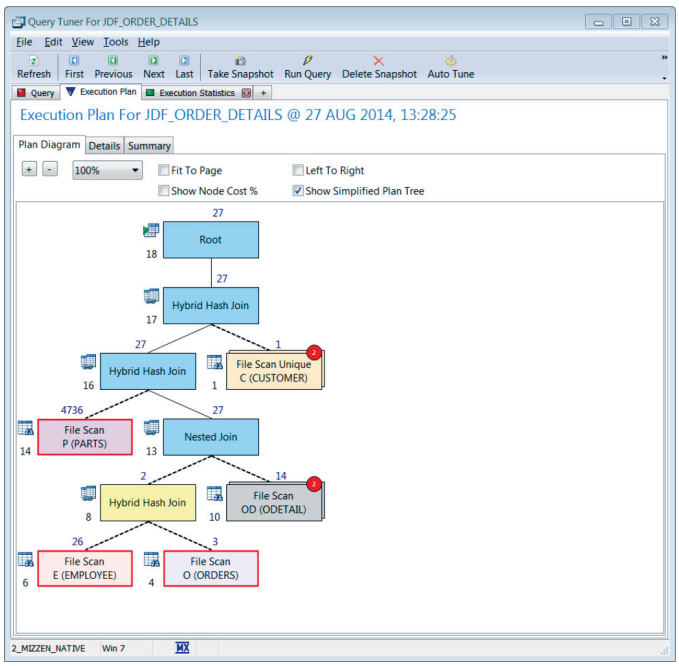

There are some tools from HPE that can help:

- The Visual Query Planner (VQP) ships with MX. It displays a plan outline in a standard Windows Tree View. This is a modest improvement; however, the tree view becomes hard to interpret as the number of operators in the plan increases.
- The Visual Query Analyzer (VQA) is available from the ATC. It is a slightly improved version of the VQP that supports a more intuitive execution plan diagram.

Although these tools from HPE help with the display of execution plans, they do not support for the full query tuning process. In particular, they don't support measuring query performance, or comparing the outcome of query tuning experiments.

There are some third-party NonStop database management solutions which are significantly more capable. One of these is SQLXPress from Merlon Software Corporation. SQLXPress is a comprehensive database management solution that includes a Visual Query Tuner.

The SQLXPress Visual Query Tuner (VQT) supports the full query tuning process, including:

- Creating query tuning "snapshots" which capture the history of your query tuning experiments in a query tuning database.
- Displaying intuitive execution plan diagrams with an easy to use drill-down feature for access to operator details. See figure 1 for an example.
- Measuring the query performance for each query tuning experiment.
- Comparing query performance against the baseline.
- Implementing changes including Update Statistics, Create Index, changing query text, and using SQL control statements.

It is beyond the scope of this article to go into further detail on understanding execution plans. A more in-depth explanation

is available in the "Reading an Execution Plan" chapter of the SQLXPress Visual Query Tuner User's Guide, which is available for free download from the Merlon website (see link at end of article).

Measuring Query Performance

In query tuning, there are two motivations for measuring performance:

- To compare the relative performance of your tuning experiments. This enables you to determine if an experiment improves performance compared to the baseline.
- To measure the amount of work being performed when executing the query. This can help identify unnecessary work, for example database records that are read, but not used.

The gold standard for monitoring performance is to use Measure. You can use Measure to collect very detailed information on all aspects of query performance. However, it is quite labor intensive to gather and analyze a complete set of Measure data for a query. This is especially true for non-trivial queries that access many tables and involve many ESP and DP2 processes. A more practical approach is to use the execution statistics provided by SQL.

SQL execution statistics include the elapsed time for compiling and executing the query, as well as per-table statistics. The per-table statistics include:

- Records read and used. The number of records read from a table and the number that were actually used.
- Number of messages and message bytes. This is a measure of message system activity, and is a good analog of IO load.
- Lock waits and escalations. This is an indicator of lock contention. The response time of the query will suffer if it encounters lock waits.

Unfortunately, neither VQP nor VQA will gather the SQL execution statistics for you. If you are using one of these tools for the execution plan analysis, you will need to use MXCI to gather the execution statistics. The easiest way to do this is to use the SET STATISTICS command. (MP users can use the SQLCI SET SESSION STATISTICS command.) Once the statistics option is turned on, the execution statistics will be displayed automatically after preparing and/or executing the query.

The SQLXPress Visual Query Tuner automatically collects and stores execution statistics as part of performing a tuning experiment.

Making Changes to Improve Performance

There are a number of techniques you can use to influence the SQL compiler's choice of execution plan:

Performing Update Statistics

The SQL compiler uses table statistics to help estimate the cost of potential execution plan candidates. (By the way, don't confuse table statistics with the execution statistics we covered above. Table statistics tell us about the distribution of data values in the columns of a table.)

Table statistics are not maintained automatically. You have to tell NonStop SQL to re-calculate them on an on-going basis. You use the UPDATE STATISTICS command for this.

Not only is it important to keep table statistics up to date, you must also generate the right statistics. In general, it is a good idea to update statistics for all key columns. If you are tuning a poorly performing query, make sure there are statistics for the columns involved in the filter conditions. For MX, be sure to include multi-column statistics when a table filter refers to more than one column.

Creating or Modifying Indexes

Adding an index is often the best way to improve query performance.

However, you need to consider the additional cost that will be incurred for INSERT, UPDATE, and DELETE (IUD) operations against the table. An insert or delete against the table will also cause an insert or delete against all of its indexes. An update will cause an index to be updated if the index includes any columns that are changed by the update.

For a table that is not too volatile, creating an index to support an important query is often worth the extra cost incurred by IUD operations.

You might consider creating an index to:

- Support a filter, avoiding a full table scan.
- Support a join, if you suspect a nested join supported by an index will perform better than a hash join.
- Support an aggregate, enabling the use of an index only scan instead of a table scan.

You might consider adding columns to an index to avoid a table scan. If an index contains all of the columns required by a query an index-only scan can be used, this is almost always more efficient than a table scan (since index rows tend to be much smaller than table rows).

Changing the Query Text

It is often possible to write a query in more than one way, yet still get the same rows in the result. Changing the way that a query is written can allow the SQL compiler to choose a better execution plan.

Here is a trivial example to illustrate the point:

Consider this query:

```
SELECT * FROM DEPT D
WHERE D.DEPTNUM IN (1000, 3000, 4000)
```

The DEPT table has a primary (cluster) key of DEPTNUM.

The compiler generates a plan that performs a full table scan on DEPT, and applies the filter predicate to each row.

If we re-write the query like this:

```
SELECT * FROM DEPT D
WHERE D.DEPTNUM IN (1000, 3000, 4000) AND
D.DEPTNUM >= 1000 AND
D.DEPTNUM <= 4000
```

This returns the same results, but the compiler generates a plan that does a subset scan that only reads the rows between DEPTNUM 1000 and 4000.

For more information, the SQLXPress Visual Query Tuner User's Guide has a tutorial chapter on re-writing queries.

Using SQL CONTROL Statements

You can use SQL CONTROL statements to influence the behavior of the SQL compiler. You can use them to directly control join order, join methods, access paths, and parallelism. There are also less direct options like specifying selectivity estimates for example.

The MX CONTROL QUERY DEFAULT (CQD) statement supports a large number of options, including a number that are useful in query tuning.

The JOIN_ORDER_BY_USER attribute allows you to control the join order used in the execution plan. The tables are joined in the order in which they appear in the FROM clause.

The HASH_JOINS, NESTED_JOINS, and MERGE_JOINS attributes can be used to enable or disable the corresponding join methods. Unfortunately, they apply to the query as a whole, and not to individual tables.

The INTERACTIVE_ACCESS and INDEX_ELIMINATION attributes can be used to encourage or discourage the use of indexes.

The ATTEMPT_ESP_PARALLELISM attribute controls the generation of parallel plans. There are many more, the MX Reference Manual has details.

The MX CONTROL QUERY SHAPE (CQS) statement provides a means to force the use of a particular execution plan. However, its use is generally discouraged. A change to the database or the SQL compiler may make the forced plan inoperable. CQS statements are hard to write, especially from scratch. The best approach is to use CQD statements (including the JON_ORDER_BY_USER option) to get a plan that is as close as possible to the desired plan. Then use the SHOWSHAPE statement to generate the CQS statement that reflects the current plan. Finally, make adjustments to the generated CQS statement.

The MP CONTROL TABLE statement provides direct control over the join order, join methods, and access paths used in the execution plan. The CONTROL EXECUTOR statement allows for the selection of parallel or serial plans.

Using CONTROL statements is a two-edged sword. They give you a way to influence the execution plan produced by the compiler, but at the same time they may prevent the compiler from finding a better plan.

Information on the use of CONTROL statements can be found in the NonStop SQL Query Guides.

Summary

The goal of query tuning is to make the query run faster. You begin by studying the execution plan. From the plan you determine why the query is performing poorly. You then decide on a plan that will improve performance and make changes to get the plan you want.

Resources

Here are some resources that you might find useful:

- SQLXPress Visual Query Tuner User's Guide. Visit www.merlon.com and follow the links to Howtos And Technical Documentation. Contains tutorial information on how to read NonStop SQL execution plans, and how to change query text to improve performance.
- SQL/MX Query Guide. Visit HP NonStop Docs and follow the links. Explains how the MX compiler works, SQL access methods, tuning tips, MX execution plan operators, and parallel plans.
- SQL/MP Query Guide. Visit HP NonStop Docs and follow the links. Covers similar ground as the MX query guide. Has an excellent description of how selectivity is calculated.
- SQL Tuning by Dan Tow, available at http://www.amazon.com/SQL-Tuning-Dan-Tow/dp/0596005733 . Describes a manual tuning methodology based on creating a query join skeleton as a graph with directed links. Not the easiest read, but it contains a lot of good information. Well worth the effort.
- SQL Performance Tuning by Peter Gulutzan and Trudy Pelzer, available at http://www.amazon.com/SQL-Performance-Tuning-Peter-Gulutzan/dp/0201791692. Covers a large number of tuning issues and describes how to deal with them for eight major databases. Does not cover NonStop SQL, but a lot of the advice is applicable.
- Joe Celko's SQL for Smarties. Available at http://www.amazon.com/Joe-Celkos-Smarties-Fourth-Edition/dp/0123820227. This is not a book about SQL tuning, but contains lots of good tips and techniques to write good queries in the first place. As the name suggests, this is not a book for beginners.

John Furlong is the development team leader for SQLXPress at Merlon Software Corp., and is responsible for the design of the Visual Query Tuner. John has over thirty years' experience working on the HPE NonStop platform. For more information about SQLXPress or Merlon, visit www.xypro.com.

Keep your hands away from the controls and nobody will get hurt!

Richard Buckle >> **CEO** >> Pyalla Technologies, LLC.

One of my lasting memories of 2015 came when we took our company command center (our RV) out onto the dunes surrounding Cape Hatteras, N.C. The MATUG event held outside Washington D.C. had wrapped up the previous day and with the potential for a long weekend before having to head to Raleigh, N.C., spending time this far out into the Atlantic seemed a lot more appealing than walking the streets of Raleigh. The road to the outer banks took us past Kitty Hawk and the Kill Devil Hills, the site of the first heavier than man flight, last century's equivalent of Armstrong's first steps on the moon.

Flight has been a big part of my family conversations and for a time I was thinking about becoming a flight engineer but unfortunately when it came time to enter the work force, the role of the flight engineer was being phased out. As for being a pilot, well, that's a whole different story and for those that know me well, I sense a collective sigh of relief that this wasn't something I seriously considered pursuing. Those stories about landing with the wheels still retracted let alone taking off with empty fuel tanks well, enough said. However, even the role of modern day pilots is under threat as increasingly, it is humans who are the weak link.

In an April 14, 2016, article in the Sydney Morning Herald magazine, Executive Style, journalist Mark Gambino wrote of Flight training 101: how to get your pilot's wings. At one point, Gambino added a cautionary note to those harboring thoughts of becoming a commercial pilot. Quoting Neel Khokhani, founder and CEO of Australian flight training school Soar Aviation, who said, "Some airlines operate on the 200-feet/200-feet rule. The pilot takes off then pushes a button when they reach 200 feet to turn on the autopilot and the computer flies the rest of the way. They're then only allowed to touch the controls when the plane is on the way down, after it descends to 200 feet." To add insult to injury, Khokhani added, "The operations manual for an airline actually restricts humans from flying a plane – you're only there for situations when the computer doesn't work."

A modern airline has such complex systems all interacting that to many outside the industry, it's probably a whole lot better if the act of flying is left to the systems themselves. Almost all situations likely to be encountered following decades of observations, can be programmed to where it's not only a case of being in a position to respond but rather, early signs of imminent danger can be immediately addressed. As far back as January, 1997, The New York Times reporter Adam Bryant wrote an article, Pilots Just Want a Little R-E-S-P-E-C-T where he noted that, "Pilots, for example, do not command the respect they once did when air travel was glamorous and only for the well-heeled." Perhaps even more relevant to an IT community, "Today, pilots have in many ways become systems managers, monitoring computers that can largely fly planes on their own. There is a joke that the cockpit of the future will have one seat for a pilot and one for a dog - the dog is there to bite the pilot's hand if he touches anything; the pilot's job is to feed the dog."

When I first began to learn the discipline of coding under the watchful eye of an IBM instructor, my fellow students and I were given tips on how to act when in the computer room. This was back when testing, late at night, meant sitting at the console of an IBM mainframe and resorting to stepping through our code, one instruction at a time. IBM mainframes had a dial you could rotate and position on "single cycle" just to be able to perform this function. As with the joke about pilots, IBM told us that in the future there would only be the operator and two dogs – to ensure there was absolutely no way the operator could touch the console. Even today, when Margo and I think about acquiring a guard dog, we always discuss the matter in terms of there being a pair of them and yes, with great affection, we would have to name them Primary and Backup.

Central to the theme of these anecdotes is the issue of scale. Complexity heaped upon complexity brings with it the recognition that few among us would ever be equipped to respond to a new situation without hesitation and do so with accuracy. The premise of the 1980s movie, War Games, was to take out of the equation missile silo operators, as even with two operators, in a time of crises they couldn't be relied upon to launch the missiles. Scale is a recognized problem leading to unimaginable complexity but so too is human frailty. Modern systems cannot afford to revert to a single-cycle mode of operation while we watch the console and go, um?

In my Back for More … column in the January / February, 2016, of The Connection, I included material from a televised keynote presentation by Martin Fink, EVP and CTO of HPE, that can be viewed on YouTube. According to Martin, HPE is making huge bets on business transforming to a hybrid infrastructure. However, Martin noted that the issue with transformation quickly becomes one of scale. As your business continues to demand faster insights from more data, Martin suggests, then clearly there will be more data and with more data there are more apps and this ends up meaning there is more infrastructure to where eventually you need more people – developers and administrators. And of course, all the while, demanding even more energy. "How many servers can one person administer?" asked Martin. Bringing IoT into the discussion, Martin then added, "How many IoT devices can one person administer?"

Clearly, where HPE is headed administrators are destined to go the way of pilots as operations ready themselves for a world of silence – once there is a wheels-up moment and the craft has cleared 200', let the machine manage the machine. It's no accident that for

such a long time operations staff, administering to the applications in support of the business, does so from a console-adorned flight deck. The imagery has been inescapable for all that time. But as we move from simply checking alerts and alarms, sensitive as we have been to the visual representation of the health of the applications (usually depicted in dials and charts in the color, "green") to where operations management of the applications themselves provides greater insight as to what will likely happen should corrective steps not be taken, we still have a long way to go before it's time to bring in the guard dogs.

Those working closely with operations management and providing application monitoring solutions today are acutely aware of the transformation under way within data centers. A number of steps have been taken, with more to occur as technology matures, but there's a clear upward trajectory, according to IR's Darc Rasmussen "where our goal is to create an operations management platform that will let you achieve more." The upward trajectory of the Prognosis we are all familiar with – the visualization of events and alarms for operations and, increasingly today, line of business managers is just the start. According to IR, when it comes to operations management down the road, it will feature better prediction and prescription properties so there aren't any surprises, which in turn will lead to greater self-healing capabilities. As IR views the world of operations management, current trends indicate the technology is headed to a fully virtualized, software-defined-everything, where the ability to become predictive, prescriptive, and self-healing will become mandatory attribute for any future operations management solution.

What will the future of data centers look like? Will there be operators in the traditional sense? Will there even be consoles? Data stream analytics, an outcome of the push for Big Data where patterns and changes in behavior are mined to better serve consumers is quickly finding a home in operations. In a recent webinar sponsored by Striim, Inc., where Striim Founder and CTO, Steve Wilkes, was joined by TDWI Research Director for Data Management, Philip Russom. When it came to use cases where ingestion (of the results from data stream analytics) played an important role, Russom highlighted operational BI, dashboards, performance management, business monitoring along with real time analytics and complex event processing (CEP). Furthermore, when a live poll was performed midway through the webinar, the results highlighted the value to business top of the list were enhancing the customer experience, followed by collecting and analyzing database change in real time and then, in third place, monitoring infrastructure, equipment or replication.

If it will eventually fall to data stream analytics solutions to target the monitoring of infrastructure and by infrastructure (I would let infrastructure become the blanket thrown over all of operations beyond just the equipment), then it only adds support for the growing popularity of the term DevOps. In other words, smaller groups devoted to rolling out new applications will in turn be the responsible for managing these applications as and when the data center operations allows them "to touch the controls when the plane is on the way down, after it descends to 200 feet!" Dramatic? Perhaps, but flying a plane and operating a system has a lot more similarities than any potential differences. After all, a plane is as complex a hybrid system as any business contemplates transforming to – yes, a hybrid that at the very least, combines the mechanics of flight with the comfort of passengers.

In many ways, the glamor (if there ever was such a consideration) of operations management is headed in the same direction as aircraft piloting. There will always be someone around, a party with the experience necessary to step in should the situation demand but for every other occasion their presence may be more to reassure everyone around them than to actually intervene every time a green dial turns red. The NonStop community has been well-served by the NonStop vendor community with a rich offering of operations management solutions and the recognition by some vendors that their path is taking them deeper into ingestion of data analytics as they become more prescriptive in their actions and even self-healing in their responses, on the same trajectory as Martin Fink views the task ahead as HPE to be continuing to pursue transformation. Very soon the question "How many servers can one person administer?" will be viewed as a trick question given that the only correct response will be "None!"

Richard Buckle is the founder and CEO of Pyalla Technologies, LLC. He has enjoyed a long association with the IT industry as a user, vendor, and more recently, as an industry commentator. Richard has over 25 years of research experience with HP's NonStop platform, including eight years working at Tandem Computers, followed by just as many years at InSession Inc. and ACI Worldwide.

Well known to the user communities of HP and IBM, Richard served as a Director of ITUG (2000- 2006), as its Chairman (2004-2005), and as the Director of Marketing of the IBM user group, SHARE, (2007-2008). Richard provides industry commentary and opinions through his community blog and you can follow him at www.itug-connection. blogspot.com, as well as through his industry association and vendor blogs, web publications and eNewsletters.

The quotes come from some of Richard's clients including HP, Integrated Research, comForte, DataExpress, Striim, Inc., InfraSoft, and OmniPayments, Inc.

Policy Based Data Management Crosses All Server Platforms, Includes All Storage Technologies

Glenn Garrahan >> Director of HPE Business >> Tributary Systems

Recall the closing scene of the original movie, Raiders of the Lost Ark; an indelible image was burnt into the minds of all theatergoers – a vast government warehouse of crates stretching on into infinity. The implication being that stored in such an environment, the rediscovered ark would prove irretrievable no matter how much energy was expended trying to retrieve the treasured antiquity. In time, any paper trail recounting its location would be surely lost……..

For data center managers a similar image must come to mind as they consider just how many files they have archived -- a situation that has only been compounded over time following decades of aggressive M&A activities. Simplistic measures involving costly separate backup solutions for proprietary server platforms such as HP NonStop, HP OpenVMS, IBM PowerSystems and IBM Mainframes no longer represent cost effective solutions. The arrival of purpose-built platforms dedicated to the provision of file backups, together with protecting information critical to the enterprise has become the preferred solution for even the most resource conscious enterprise.

Most companies have cut staff while increasing their dependence on IT, requiring fewer employees to be more productive and data management processes more efficient. As a result, even short interruptions to data access are intolerable. Business decision makers and IT managers understand the need to protect data in secondary or even tertiary disaster recovery sites while continuing to ensure its immediate availability.

No longer is it good enough for production applications to have their own, dedicated data protection and file backup solutions – the world of enterprise IT is far too complex for such simplistic and anachronistic arrangements. Instead, data management solutions addressing all the needs of the entire data center are needed to assure enterprises that their data is protected and their business-critical files backed up. It is the responsibility of data center managers to develop strategies that ensure enterprises are never exposed to information loss for any reason and in so doing, execute on such strategies at an affordable price. When competitive pressures are felt enterprise wide, it pains CIOs to see resources squandered and efficiencies compromised and left unabated, poorly implemented solutions expose such enterprises to potentially damaging losses.

At the very heart of IT responsibilities is the mandate that IT can dig deep into past files and recover data pertinent to any requirement an enterprise may face. The media selected for this purpose had to be both inexpensive and yet reliable and most importantly, open and supportive of industry standards. However, the problem for many IT groups is that just as selection of applications has led to proprietary processors becoming technology islands, so have their backup solutions. The need to consolidate and embrace virtualization is now every bit as important as breaking down these legacy application silos.

Now individual storage vendors will tout their own approach as best and the new normal in data protection and consolidation. However, often these are additional servers (usually employing either a Linux or Windows OS) dropped into the existing infrastructure to address only specific proprietary host platforms, when the optimal solution is a cross-platform approach. The challenge is finding a sustainable way to build and maintain a policy-based data protection and converged infrastructure strategy that encompasses all host platforms and storage technologies, applying appropriate services to multiple data sets based on the data's criticality. In other words, "Intelligent Converged Backup Solution Methodology."

Not surprisingly, many vendor-guided discussions of backup solution methodology quickly devolve into an either/or debate over tape versus disk and specific proprietary platform support. Storage array manufacturers promote disk-to-disk mirroring (between two or more arrays on the same raised floor) or WAN-based replication (copying data between disk arrays in two geographically separate locations across a telecommunications facility). Tape technology vendors note that as much as 87 percent of the world's data is still protected using their magnetic media, even in the "Cloud". Then on top of the tape versus disk argument there's the issue of supporting proprietary versus open systems.

In actuality, both disk and tape are necessary and required. What's really needed is a solution methodology that incorporates both disk and tape backup, with availability to any host platform in a customer's IT environment. In other words a software defined, policy based strategy for all-inclusive data management.

Needed: Intelligent Backup Data Analysis

The fact is all data is not the same. Data derives its importance – its criticality to the business and its priority for access following an interruption event – from the business process it supports. That's why intelligent data management planning must be preceded by the due diligence of business process analysis. Determination of which business processes are mission critical and which are merely important is mandatory. A physical mapping of business processes to their data (e.g., data produced and used by that process) must be performed, and subsequently, of data to its actual location in the infrastructure.

This is the heavy lifting of data management, and it needs to be done for several reasons:

1. Because all data isn't the same in terms of criticality or priority of restore, there's no one-size-fits-all data protection strategy. In most organizations, data protection involves a mixture of services used in different combinations to meet the protection and recovery requirements of different business processes and their data.

2. The diversity of threats to data – including bit errors occurring close to the data itself, an array failure, or a power outage with a broad geographical footprint, occurring outside the disk drive – might require different protective services. "Defense in depth" is a term often associated with this reimagining of policy-based data protection.
3. The practical issue of cost is a factor: one-size-fits-all strategies for data management tend to not be very cost effective. Backing up mission-critical data to tape might create problems with the recovery requirements of always-on applications, because data restore might take longer than the time-to-data requirements of such applications. Conversely, using an expensive WAN service to replicate the data of a low-priority application or business process could well be regarded as overkill and a waste of money.

So, getting more granular in the assignment of management, protection and recovery services to data assets – creating policies for different data sets based on a thorough analysis of requirements – is the best way to rationalize expense while streamlining and optimizing data backup.

Building an Integrated Data Management Strategy

Once the business needs of data are understood and mapped, the challenge is building and maintaining a data management strategy that includes multiple services targeted at multiple data sets.

One approach would be to centralize service delivery using storage virtualization, which abstracts value-add services (like mirroring, replication, snapshots, etc.) away from storage arrays and centralizes them in the storage virtualization software "uber-controller." But to the best of our knowledge no universal storage virtualization product works with the full range of enterprise host and storage sub-systems.

In the absence of such a centralized strategy, another location is needed where data tends to aggregate, hence providing the opportunity to apply policy based data management services. One idea is to use a Virtual Tape Library (VTL), which has come into widespread use over the past decade.

VTL technology evolved from a location where backup jobs were stacked until sufficient data existed to fully fill a tape cartridge (first generation) to an appliance offering virtual tape devices to expedite backup (second generation). Today, VTLs have advanced to a location where 30, 60, or 90 days' worth of data is stored for fast restore in the case of an individual file corruption event (one of the more common types of data disasters). VTLs have also been enhanced with additional storage services, including VTL-to-VTL replication across WANs, de-duplication or compression services, and even as the place to apply encryption to data.

The VTL is already a location where much of the corporate data "hangs out" and where various data management services can be readily applied. It makes sense to make this platform a "storage services director," a location where data protection, data management and archiving services could be applied to data assets per pre-established policies.

This federated approach, in which a "storage director" serves as the cornerstone of policy-based data management services delivery, is gaining mindshare among users. Such storage directors must be able to be clustered to facilitate the handling of large quantities of data via clustered nodes, thus ensuring adequate connectivity for

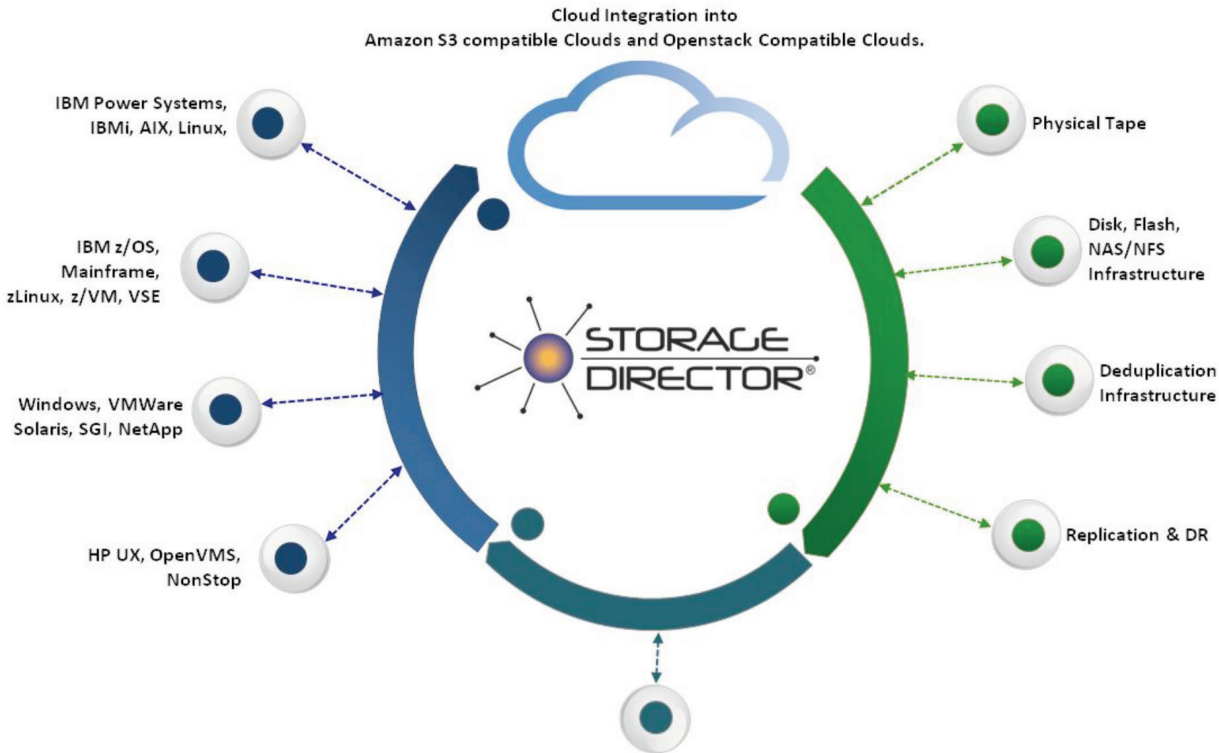

All Backup Application Support:Veritas NetBackup, CommVault Simpana, Tivoli Storage Manager, Veritas Backup Exec, CA ARCserve, Veeam, HP Data Protector, etc.

Protect your most sensitive data from damaging breaches

Make your most sensitive data useless to hackers. Encrypt and tokenize it—at the moment it's created and throughout its lifecycle. HPE SecureData protects your data at rest, in motion, and in use. On-site, in the cloud, on mobile endpoints, and in big data platforms. Our continuous data-centric security keeps your business out of the headlines. It helps enable PCI compliance and audit scope reduction, and de-identifies data to protect personally identifiable information, protected health information, and intellectual property. HPE Security—Data Security format-preserving encryption, stateless key management and tokenization solutions keep you ahead of hackers.

Learn more at:
voltage.com
hpe.com/go/DataSecurity

Visit us at Booth #114

Hewlett Packard Enterprise

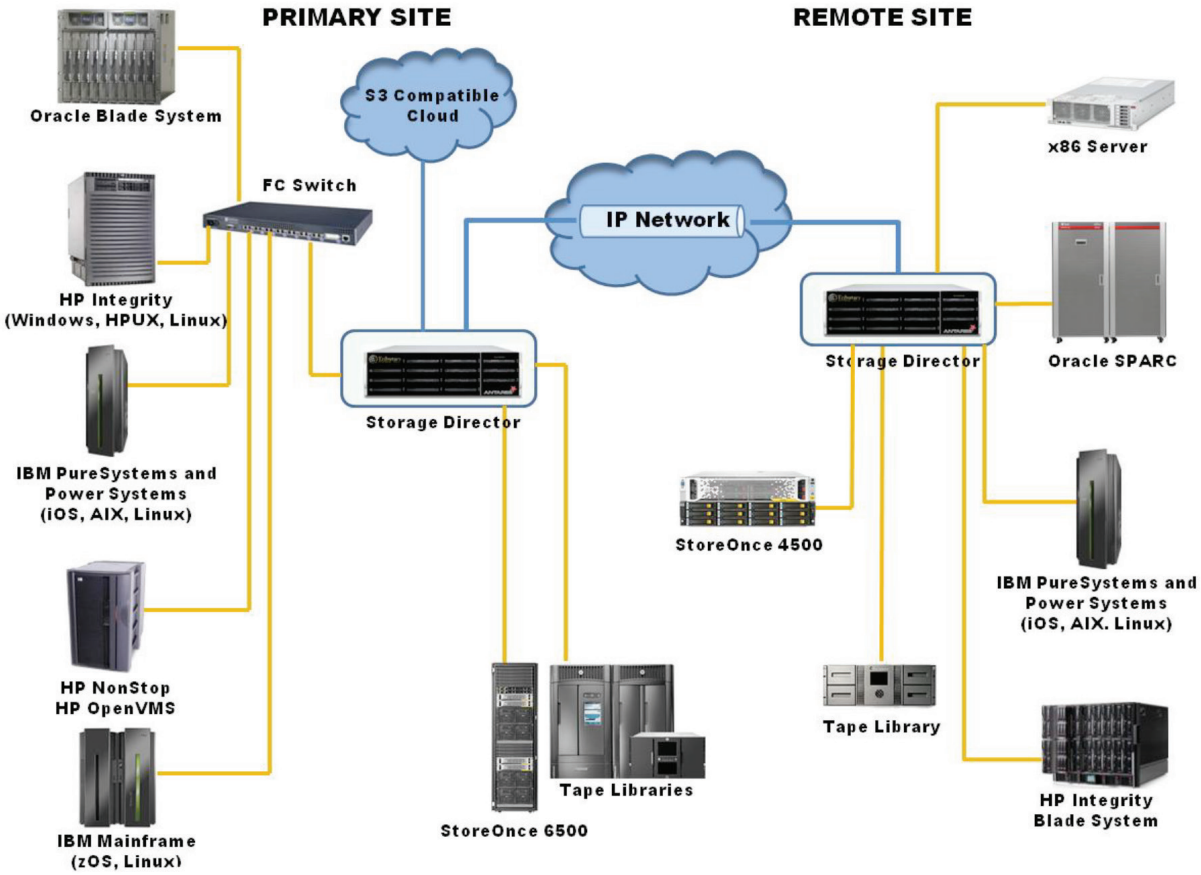

both data sources and converged data storage infrastructure while streamlining the management of both the directors themselves and the policies they implement.

The storage director itself may be software that can be implemented on any hardware server platform or may be a hardware appliance preconfigured with software for ease of deployment. More importantly, this storage director must be capable of being dropped into any complex of servers and storage devices, supporting all connectivity protocols, to minimize disruptions. With those capabilities, organizations will have moved closer to the ideal of managing data assets with a converged, sustainable, defense-in-depth strategy.

Storage Director 5.0 by Tributary Systems

One company that has taken this approach is Tributary Systems; whose Storage Director® 5.0 product is a software defined policy-based data management solution that connects "any host to any storage technology". Storage Director applies appropriate protective, management and archival services to data, while simultaneously converging the backup infrastructure. In fact, Tributary Systems is an HPE Converged Infrastructure Ready Solutions Partner.

With Tributary Systems' Storage Director, enterprises can apply data policies down to individual data volumes on host platforms (such as HPE NonStop), based on business criteria and importance to business resiliency and restoration. This is intelligent data management!

In enterprises with multiple host platforms – HPE NonStop NB, NS and now NonStop X servers, HPE Open VMS, Windows and VMware running HPE Data Protector, IBM zOS mainframes, IBM AS/400s iOS (now IBM PowerSystems), among others – Storage Director enables sharing storage technologies otherwise dedicated to each host platform. Such storage technologies can include existing enterprise storage disk, HPE StoreOnce, EMC Data Domain, and Quantum DXI data de-duplication devices, physical tape, Amazon S3 and Openstack compatible Cloud, or any combination of storage technologies dictated by individual data management needs. Such a converged approach improves storage performance, enables consolidation, and can lead to real savings of time, power and most importantly staff.

Conclusion

The dependence of business on automated systems to support critical processes has never been greater than it is today. Even a short-term interruption in access to data required by applications and by decision makers can have disastrous consequences. This situation underscores the requirement to develop, implement and maintain a comprehensive, cohesive, cost-effective policy-based data management strategy. Such a strategy must be guided by the recovery requirements and timeframes of business processes themselves and must also encompass the many data management tools and techniques available.

Without a doubt it's time to move beyond platform-specific solutions, and instead consider an intelligent, converged approach to data management.

Glenn Garrahan is Director of HP Business for Tributary Systems, Inc. Prior to this, Glenn spent 17 years as a Product and Program Manager with HP NonStop

>>Native Tables >>in NonStop SQL/MX Part 2

Frans Jongma >> **Master Technologist** >> NonStop Advanced Technology Center

Introduction

This is part two of a two-part article about the features of Native tables in NonStop SQL/MX. Part one was included in the March-April issue 2016 of this magazine.

DML variants that are allowed only on native tables

In the earlier releases, all SQL/MX DML could be executed on SQL/MP tables by the SQL/MX engine. With releases newer than 3.0, some new DML statements have been introduced that will only work on native tables. New, native-only DML features are: updatable primary keys, self-referencing updates and multi-commit delete.

Updating the primary key value

With the release of SQL/MX 3.2, updates on the primary keys are allowed. Note however, that this update is implemented as a delete of the row(s) followed by an insert. The reason is simple: as a result of the update of the primary key the row may have to move from one partition to another.

Self-referencing updates

Self-referencing updates are updating a table that is read within the same statement using a subquery.

Imagine a gaming application where you want existing rows to be copied to new rows:

```
>>insert into tx1 select identifier + 1000,
user_name, 'New Game' from tx1;
--- 500 row(s) inserted.
```

This example shows how to clean up rows that are below an average.

```
>>delete from tx1 where u < (select avg(u)
from tx1);
--- 100 row(s) deleted.
```

SQL/MX assures that rows are only processed once. This avoids getting into a loop as is the risk with the INSERT statement where new rows with incremented identifiers are inserted at the end of the table.

Multi-commit delete

DELETE operations on very large tables can be very long running, locking many rows and may lead to lock escalations or even worse, a TMF timeout that causes the delete to be backed out. The Multi-commit option for the DELETE statement breaks up the operation into multiple smaller delete transactions. These deletes will be performed by separate Executor Server Processes (ESP), one for each partition. Currently, the multi-commit option does not allow host variables in the WHERE clause. It is therefore more useful in a dynamic SQL environment.

```
>>delete with multi commit every 200 rows
from tx1 where u < 500;
--- 4 row(s) deleted.
```

References

SQL/MX reference manual: Chapter 2, DELETE Statement.

Data types

NonStop SQL/MX native tables support floating point data in IEEE format, which differs from the Tandem-float format that is used in SQL/MP tables. Character Large Objects (CLOBS) and Binary Large Objects (BLOBS) are supported by the JDBC drivers, new data types are supported when used in native tables.

Floating point data types

When SQL/MP tables are migrated to SQL/MX native tables, the floating point data columns need to be examined for compatibility. SQL/MP uses a different format to implement floating point data than IEEE which SQL/MX uses. For example, there is no equivalent for a Tandem REAL in the IEEE data type which preserves the precision and exponent. The IEEE REAL uses less storage but supports a smaller exponent size. Alternatively, for more exponent and precision, the column can be defined as FLOAT or DOUBLE PRECISION.

Extended numeric precision

Columns defined as NUMERIC (precision, scale) can have a scale of maximum 128 digits, which is an increase from 18 in older releases.

Unicode data type

SQL/MX native tables support ISO88591 for single-byte characters and UCS2 (Unicode). Unicode is a universal encoded character set that lets you store data from any language using a single character set. NonStop SQL/MX uses the UTF-16BE (16 bits, big endian) encoding for UCS2. The full range of UTF-16

characters is allowed, but surrogate pairs are not recognized as characters. (Surrogate pairs are used to expand the code space of UTF-16 beyond the UCS-2 character encoding.)

Large Objects (CLOBs and BLOBs)

Native tables provide support for large binary and character objects, referred to as BLOBs and CLOBS. Because the largest supported row size is approximately 32KB, an object that is larger than this maximum cannot be stored in a single row. The JDBC drivers support these objects by inserting the objects into a dedicated table in as many 32K chunks as necessary to store the object. The driver takes care of assigning a key value for the object and to store the data into the table.

References

SQL/MX reference manual: Chapter 6, Language Elements: Data Types.

JDBC Type 2 Driver Programmer's Reference: Chapter 4, Working with BLOB and CLOB Data

NonStop JDBC Type 4 Driver Programmer's reference: Chapter 5, Working with BLOB and CLOB Data

Other functionality

This section describes the miscellaneous features of SQL/MX native tables. For example, the SQL/MX metadata is stored in SQL/MX native tables; the FASTCOPY utility can efficiently copy native tables and indexes and the GUARDIAN FUP program can show information about native tables using their native (ANSI) names. Backup and restore of native tables must be done by Backup/Restore2 instead of Backup/Restore.

Metadata tables

Users of SQL/MP with the SQL/MX engine use the SQL/MX metadata only to a small extent. This metadata is often only used to map SQL/MP tables to SQL/MX aliases to allow the use of full ANSI names to access the data stored in the MP tables. When native tables are used, all the information about them is stored in highly normalized SQL tables. The layout of the tables is described in Chapter 10 of the SQL/MX reference manual. The NonStop Advanced Technology Center provides a utility, mxschema , that creates a set of utility views on the metadata that allow a DBA to view information about a objects in a schema using simple SQL queries instead of writing complex joins. The utility is downloadable from the HPE Software Depot free of charge, but note that it is provided "as-is", without warranty of any kind.

Distributed databases

Native tables are accessed via the catalog they are registered in. When tables are distributed across multiple NonStop servers, the catalog must be registered on each of the nodes using the REGISTER CATALOG command to make the catalog visible to these nodes. If the catalog is not visible, a program cannot get access to the data. This is a different approach from SQL/MP tables where remote data could be accessed by their file names such as \NODE.$DATA.SQL.TBL. SQL/MX tables will always be called by their catalog.schema.table name regardless of the system they are located.

A user must be granted access via GRANT/REVOKE, but in order to access remote data, GUARDIAN REMOTEPASSWORDS must be set up, just as is required to access SQL/MP remote data.

Stored Procedures

Stored procedures can access native tables as well as SQL/MP tables using the same API and the same (JDBC T2) driver. However, the SPJ must be defined in the SQL/MX catalog.

Views

SQL/MX views can only refer to SQL/MX native tables. The SQL/MX view text is not limited to the 3,000 characters of SQL/MP.

Utilities

Many new utilities exist for SQL/MX and this article will only touch upon a few of them that are important to know about.

Fastcopy

Fastcopy is a command that is similar to SQL/MP LOAD or COPY. It can copy data from one large table to another faster than a normal insert/select. The most important things to know about fastcopy are:

- Fastcopy performs the functionality of INSERT INTO <target_table> (*) SELECT * FROM <source_table>. This means that the source and target tables must be similar, have the same number of columns, compatible data types and so on.
- The target table will be made non-audited for the duration of the command, to allow fast inserts.
- Due to the initialization, fastcopy of small tables will take longer than a normal insert-select.
- The command can be executed by user programs using dynamic SQL.
- Indexes, when defined on the target and online, are copied in a second step the same way the table was copied, in the order of the alternate key. This is more efficient than to create the index separately.
- The source table will be accessible for read-only operation while fastcopy executes.
- The syntax allows for separate commands to copy a table and index. Fastcopy will detect if any updates on the source have been applied and will not release the targets unless it completed successfully.
- System generated key values, such as SYSKEYS and IDENTITY columns are copied to the target.
- The target table and indexes will be compacted as if a reload had been performed.

FUP

Fup has been enhanced over the years to provide some support for SQL/MX native tables. However, SQL/MP users may not be fully aware of the power of the ANSI name support in FUP. The next paragraphs highlight some features of the INFO and LISTLOCKS commands.

Enhanced INFO command

FUP INFO requested on a table by the ANSI name shows information on all the partitions of the object instead of just a single Guardian file. The next example shows the information of the ORDER_EXAMPLE table. It has three partitions, one on $DATA02 and two on $DATA01. A single INFO command shows the information of all partitions.

```
-info 'table frans.perf.order_example'
```

NonStop Technical Boot Camp

SAVE THE DATE

November 13-16, 2016
Fairmont San Jose Hotel | San Jose, CA

- 50+ Technical Breakout Sessions
- Executive Keynotes, Roadmaps and MORE!
- Partner Pavlion and Partner Networking Events
- Downtown San Jose Venue

REGISTER ONLINE TODAY!
www.connect-community.org/nonstop-technical-boot-camp

```
CODE       EOF    LAST MODIF    OWNER RWEP
TYPE       REC BL
$DATA01.ZSDFJ001
DK4BGG00          550A+  12288  7:25 211,96
*SQL  XPK Ta    28  4
VT1BGG00          550A+  0 4Sep2013 6:44 211,96
*SQL  PK Ta    28  4

$DATA02.ZSDFJ001
L54BGG00          550A+  12288  7:32 211,96
*SQL  XPK Ta    28  4
```

Enhanced LISTLOCKS command

This example shows the locks placed on the ORDER_DETAIL table when a transaction has deleted one row, but has not committed yet. The table has three partitions and every partition has an "Intent-lock" placed on it. The actual row lock is in the partition on $DATA02.

```
-listlocks  'table frans.perf.order_example'

$DATA01.ZSDFJ001.DK4BGG00
ANSI NAME FRANS.PERF.ORDER_EXAMPLE
 LOCK      REQUESTER            KEY
 TYPE  STATE  ID         LEN  KEY/RECORD
 ADDRESS
 F     GI     \NSKIT10(2).1.168356

$DATA01.ZSDFJ001.VT1BGG00
ANSI NAME FRANS.PERF.ORDER_EXAMPLE
 LOCK      REQUESTER            KEY
 TYPE  STATE  ID         LEN  KEY/RECORD
 ADDRESS
 F     GI     \NSKIT10(2).1.168356

$DATA02.ZSDFJ001.L54BGG00
ANSI NAME FRANS.PERF.ORDER_EXAMPLE
 LOCK      REQUESTER            KEY
 TYPE  STATE  ID         LEN  KEY/RECORD
 ADDRESS
 R     G      \NSKIT10(2).1.168356      16
 ?0 ?0 ?0 ?4 ?0 ?1 "REG_50
"
 R     G      \NSKIT10(2).1.168356      16
 ?255 ?255 ?255 ?255 ?255
?255 ?255 ?255 ?255 ?255 ?255 ?255 ?255 ?255
?255 ?255
  F     GI     \NSKIT10(2).1.168356
 -
```

The LISTLOCKS command can also be used on the schema level by specifying the SCHEMA keyword instead TABLE.

Warning: The amount of locks displayed may be high on production systems.

Backup and Restore using BR2

SQL/MX objects must be backed up and restored using new utilities. Backup/Restore2 is used to backup and restore SQL/MX native tables and OSS files to tape.

Customers may already use PAK and UNPAK utilities to backup and restore Guardian objects to disk. PAK2 and UNPAK2 are available but not yet distributed the same way as PAK. PAK2/UNPAK2 are best executed from the OSS environment and they use the OSS file system, however, a Guardian non-native (code 100) version exists.

References

SQL/MX reference manual: Chapter 2, FASTCOPY

SQL/MX reference manual: Chapter 10, Metadata Tables

SQL/MX Guide to Stored Procedures in Java

Backup and Restore 2 Manual

Concepts of NonStop SQL/MX, Introduction to SQL/MX Metadata: HP Document ID: 4AA3-6539ENW

Concepts of NonStop SQL/MX, Introduction to SQL/MX Stored Procedures: HP document ID: 4AA4-9428ENW

PAK2/UNPAK2 information: http://193.65.99.19/kku/pak/pak4.html

Conclusion

Many SQL/MX customers are happy with the functionality that the SQL/MX engine provides them accessing their SQL/MP data. However, today's applications often require more functionality that only native tables provide. This article, along with the first part that was published in the March – April 2016 issue, has presented an overview of the features that native tables provide, and the list may be a lot longer than most SQL/MP customers expect. I did not plan to write this many pages worth of differences! Today's applications maintain more data, which results in storing larger rows; they make use of "meaningless" keys, generated by the system as –for example- sequences and they assume the DBMS to maintain Referential Integrity. All of this is supported by the native tables of SQL/MX. SQL/MX security is enhanced by the introduction of the Security Administrator role: the SA can manage the security of database objects without having access to the data itself. Note that SQL/MX native tables require different backup and restore procedures than SQL/MP data. Backups of native tables require the Backup/Restore2 product and can be backed up to disk using the pak2/unpak2 utilities. These two utilities will be included with the software in the upcoming SQL/MX release.

Frans Jongma is a Master Technologist for the NonStop Advanced Technology Center (ATC) and is based in Europe in The Netherlands. Frans has worked in several consulting positions for the NonStop Enterprise Division since 1989. His main areas of expertise are: NonStop SQL (MP as well as MX), application design, performance analysis and high-availability. Prior to joining Tandem, Frans has worked on the design and implementation of database management systems and developer productivity tools for UNIX and proprietary systems. Over the years he has been advocating the use of new technologies that operate in the HPE Open Systems Services (OSS) environment, such as NonStop Server for Java and NonStop SQL/MX software.

The Dawn of Fault-Tolerant Computing

Dr. Bill Highleyman >> Managing Editor >> Availability Digest

In 1980, I published a four-part series in Computerworld entitled "Survivable Systems." The articles described the state-of-the-art fault-tolerant systems at the time. The need for systems that never (at least, hardly ever) failed was just being recognized. Several companies jumped in with their own versions of fault-tolerant systems, including Tandem, Stratus, Synapse, Auragen, August, NoHalt, Parallel Computers, and Tolerant Systems.

A lot has changed over the 36 years. Systems have become more "open," with Linux-like operating systems and x86-based hardware architectures. However, what hasn't changed is the need for systems that never fail. Applications that were hardly in use in the 1980s now are becoming mission-critical. The use of internet email is a perfect example. With the advent of social media, systems promoted as 24x7 can't risk a failure. As soon as a system is under distress, the Twitter universe explodes with complaints and comments, often causing irreparable harm to a company's reputation for reliability.

Some early products are still in use, for instance, Tandem and Stratus systems. Others have been incorporated into newer products. Still others simply have disappeared. In this article, we visit the dawn of fault-tolerant computers and the various architectures that were being promoted as such at the time.

Tandem Computers, Inc.

The Tandem computer was the granddaddy of fault-tolerant systems. Tandem's first system was delivered in 1976. Forty years later, its original architecture remains the dominant fault-tolerant

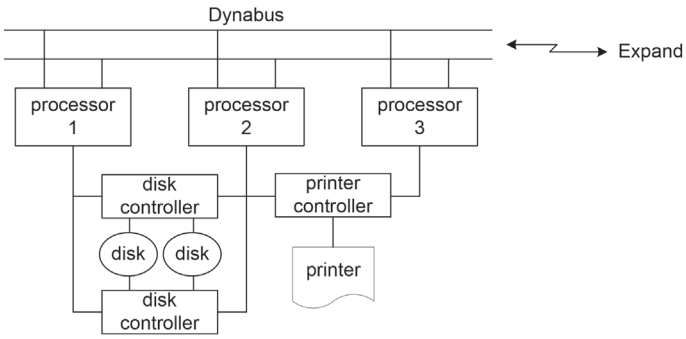

Tandem

technology. Then and now, a Tandem system was a loosely coupled multiprocessor system that contained anywhere from two to sixteen independent processors in a node. Up to 255 nodes could be included in a single network, linked via Tandem's Expand communication network.

The processors in a node were linked via a duplexed, interprocessor messaging bus called the Dynabus, capable of a 26 megabyte/second data rate.

All device controllers were dual-ported so that there was always a path to a device even if a processor failed. All critical processes ran as process pairs in two different processors. One process was the primary process, and one was the backup process. The primary process kept its backup process synchronized via checkpointing messages. Should the primary process fail (presumably due to a processor failure), the backup process took over and continued processing with no apparent interruption to the user. (Tandem's later inclusion of the Pathway process monitor eliminated the need for application programmers to write checkpointed process pairs.)

With Tandem's second release of its product, the Tandem NS2, each processor could be configured with two megabytes of memory. Each mirrored disk pair could provide 128 megabytes of memory (yes, that's megabytes, not gigabytes).

Tandem was acquired by Compaq in 1997, which then was acquired by HP in 2002. Tandem computers are now known as HPE NonStop computers. (HPE is HP Enterprise, one of the two companies that resulted from the split in 2015 of Hewlett Packard into HP, Inc., which sells HP personal computers and printers, and HPE, which markets HP server, storage, and networking systems.)

Stratus Computer, Inc.

Founded in 1980, Stratus Computer, Inc. (now Stratus Technologies, Inc.[1]) marketed a system that was similar to the Tandem system in that it was an expandable multiprocessor system. However, the similarity ended there. Stratus achieved with closely-coupled hardware what Tandem achieved with loosely-coupled software.

Each Stratus processing module comprised four Motorola 68000 microprocessors. Introduced in 1967, the M68000 architecture is still in use today, almost four decades later.

A Stratus CPU board contained two such microprocessors running in lockstep. The outputs of the two microprocessors were continuously compared. As long as they behaved identically, their outputs drove a pair of high-speed buses that communicated with other modules in the system. However, should a difference be detected in the outputs of the

[1] In 1998, Stratus Computer, Inc. was purchased by Ascend Communications for its communication products. The enterprise server portion of the business was spun off to the original founders in 1999. The new company was named Stratus Technologies, Inc.

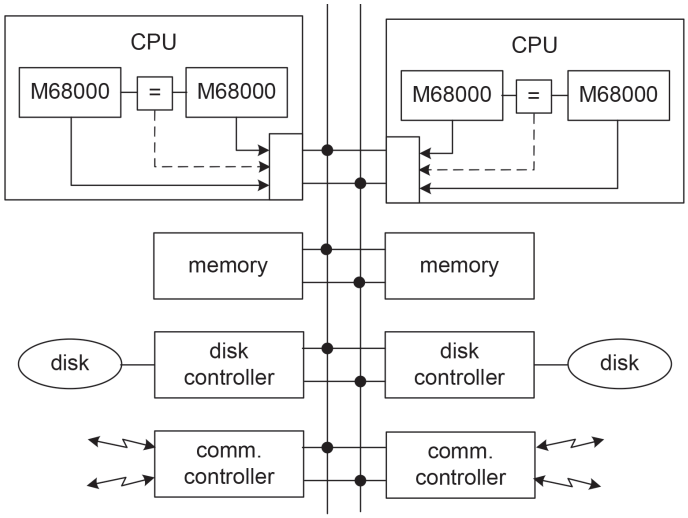

Stratus

two microprocessors, the board was shut down immediately. Thus, the CPU would not propagate any bad result.

To provide fault tolerance, two CPU boards drove the pair of buses. As long as both boards functioned properly, the buses were driven with the same signals from each board. However, if a fault was detected in one of the boards, the board would be shut down. The surviving board continued to provide the processing functions for the system.

This philosophy of hardware self-checking and redundancy was carried throughout the rest of the system via dual memories, dual disk controllers, and dual communication controllers. Multiple processing modules could be linked together via a simplexed or duplexed 2.8 megabyte/second interprocessor bus called the StrataLink. Each Stratus processor could be configured with four megabytes of memory and sixty megabytes of mirrored disk.

Over the years, Stratus' lockstep technology evolved to become its ftServer. Stratus also introduced a software solution, Avance, which used two general-purpose x86 servers that were kept synchronized with synchronous replication. Avance was replaced later by Marathon's everRun system when Stratus acquired Marathon in 2012.

In addition to these systems, in the early 1980s Stratus introduced the VOS operating system running on its high-performance Continuum hardware. VOS was a Multics-like operating system, and Continuum fault tolerance was achieved with dual processors running in lockstep. Stratus since has ported VOS to its ftServers. Interestingly, my company, The Sombers Group, was contracted by Tandem in the early 1980s to compare the performance of VOS with the Tandem TXP system. Their performances were substantially identical. However, Stratus' focus then and now remains within the market for smaller x86-based solutions.

Synapse Computer Corporation

Synapse Computer Corporation took an approach strikingly different from that taken by Tandem and Stratus. Rather than providing virtually instant recovery from a failure, Synapse took the view that a short delay in recovery was acceptable provided the database was not corrupted. Synapse argued that it was perfectly acceptable to have the users wait a minute or two while the system recovered and then to require the users to reenter the transactions

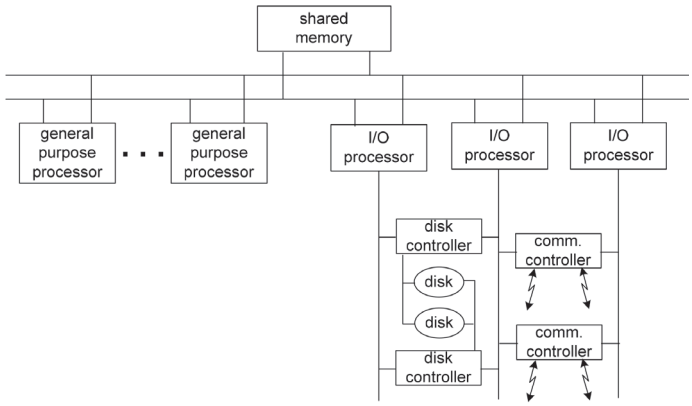

Synapse

that were in progress at the time of failure. Back then, this was called a "mission critical solution." Can you imagine trying to sell it in 2016?

The Synapse architecture was a closely coupled configuration. Up to 28 general-purpose or I/O processors could be connected to a dual 32 megabyte/second bus. The processors communicated via a shared memory also connected to the dual buses. Thus, interprocessor communication was much faster than for the Tandem and Stratus systems, both of which relied on interprocess messaging.

Also, processes were not assigned to processors. Rather, all processors shared a common task queue. When a processor became idle, it began processing the next task on the queue. A highly efficient caching mechanism allowed processors to access data anywhere, even if it were in the cache of another processor.

The Synapse system was certainly not fault-tolerant in the sense of Tandem and Stratus. If a memory module failed, every executing process in the system could be affected. If a general-purpose processor failed, every process running in that processor failed.

Synapse depended upon a transaction model to recover from failure. Transactions were checkpointed as each new screen was presented to a user. If a failure occurred, the entire system was brought to a halt and was reconfigured around the failure. The database was recovered to ensure its consistency. Database recovery was accomplished via a History Log, which contained all before and after images of data changes made to the disks. Transactions then were restarted from their last checkpoints.

A Synapse system used Motorola 68000 microprocessors. It could be configured with six megabytes of memory and 300 megabytes of disk storage.

The closely coupled architecture of Synapse systems never was accepted as a fault-tolerant solution, and Synapse systems no longer are sold.

Auragen Systems Corp.

Auragen's fault-tolerant offering was in many respects a hybrid combination of the loosely coupled architecture used by Tandem and the closely coupled architecture used by Stratus. It comprised from 2 to 32 processor clusters connected to a dual-system bus, each with a 16 megabyte/second capacity. Devices connected to the processor clusters via dual-ported controllers so that there was an access path to every device even in the event of a processor failure.

A processor cluster was a closely coupled system comprising an executive processor and two work processors, all of which were

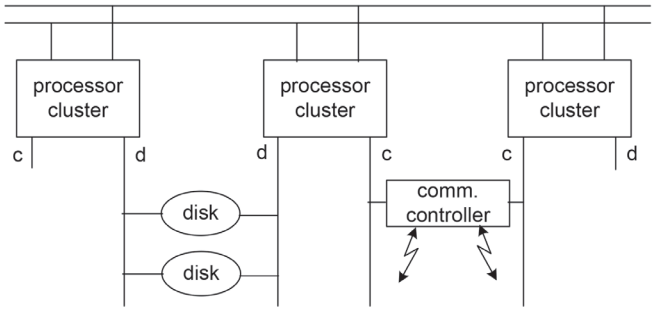

System Architecture

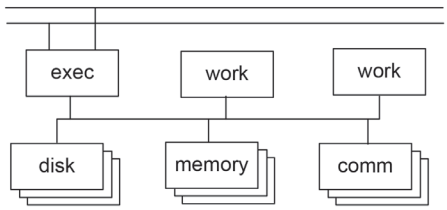

Processor Cluster

Auragen

Motorola 68000 microprocessors. The processors in a processor cluster were interconnected with shared memory and device controllers via a 20 megabyte/second bus. The executive processor was responsible for all operating-system functions. The work processors were independent from each other and ran the applications.

The processor clusters were not fault-tolerant. Rather, fault tolerance was achieved in a manner similar to Tandem systems. Each process had a backup process running in another processor cluster, and the backup process was kept synchronized with the primary process via checkpointing. However, checkpointing was taken care of automatically by the processor clusters. All communication between processes was via interprocess messages. All interprocess messages directed to the primary process also were queued to the backup process. Periodically, the backup process would be synchronized with the primary process; and its queue of messages would be erased.

Should the active process fail, the backup process replayed all of the messages in its queue. During the replay, outgoing interprocess messages were inhibited to prevent duplicates. Following the completion of the interprocess-message replay, the backup process was ready to take over processing where the active process left off.

Auragen was involuntarily dissolved in 1985.

August Systems

The August system was unique in that it provided fault tolerance for both digital and analog signals. It comprised a loosely coupled, triplexed voting system using triple modular redundant (TMR) technology. Three independent paths were provided through the system, and their results were compared. If one result was different from the other results, it was discarded. The common result was passed on as the correct output.

The August system comprised three 8086-based Control Computer Modules (CCMs) that communicated with each other over read-only interprocessor buses. Each CCM received inputs from the Processor Interface Module (PIM) and made its own calculations. The results were compared via the interprocessor buses. If a CCM found itself outvoted, it adjusted its results to comply with the other two modules. If a CCM found itself consistently outvoted, it declared itself out of service.

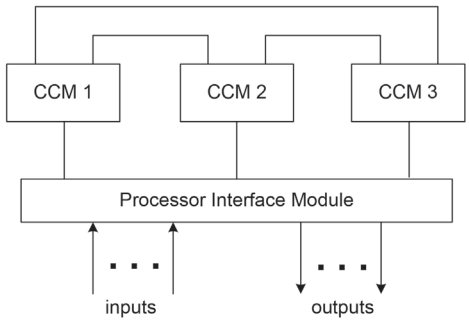

System Architecture

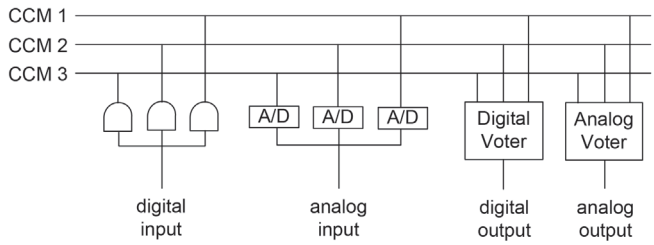

Process Interface Module

August

Each CCM returned its response to the PIM, which itself voted on the results and returned the majority result to the outside world. Analog inputs were digitized via analog-to-digital (A/D) converters. For analog output signals, the digitized values were reconverted to analog values; and the median value was returned as the result.

The PIM provided input and output redundancy for digital and analog signals. Input redundancy was provided by distributing input signals to all three CCMs. The output circuits were fault-tolerant. If any component opened or shorted, the voter still functioned properly.

Each CCM could be configured with 128 kilobytes of memory.

In 1997, August Systems was purchased by ABB, a Zurich, Switzerland-based company operating mainly in robotics and the power and automation technology areas. The August system is now known as the ABB Triguard TMR product. It has over 1,000 systems installed worldwide with over 10 million operational hours.

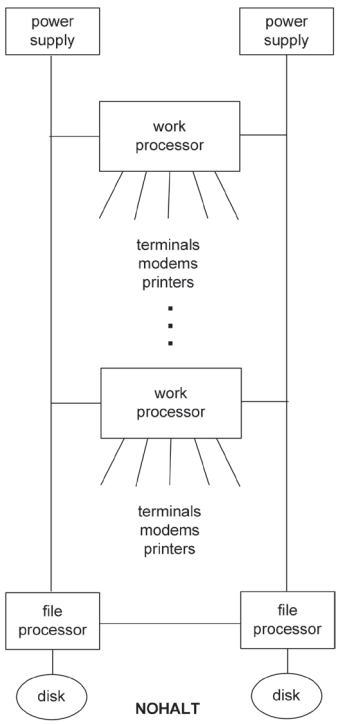

NoHalt Computers, Inc.

NoHalt Computers aimed at the low end of the market. Its system comprised a mirrored database with up to sixteen 8-bit Zilog Z80 work processors interconnected by a pair of 1.25 megabyte/sec. interprocessor buses.

Mirrored files were implemented via a pair of file processors that were independently powered and that managed their own disk units. They communicated with the

work processors via the interprocessor buses.

Each work processor could support up to four peripheral devices such as terminals, modems, or printers. However, the work processors were not multitasking. They each could perform only one task at a time.

In the event of a file-processor failure, the surviving processor continued to support the system. In the event of a work-processor failure, all peripherals connected to that processor were out of service.

Each file or work processor could be configured with 64K of memory. The system could be configured with 140 megabytes of mirrored disk.

NoHalt Computers was acquired by TPC Logistics Services, Inc., and the NoHalt product was renamed the Reliant fault-tolerant computer system. In 1984, TPC announced the Reliant product had been enhanced to support 16-bit 8086 microprocessors.

Parallel Computers, Inc.

Based on Motorola 68000 microprocessors, the Parallel Computer system could connect up to five parallel processing units (PPUs) to a simplexed interprocessor bus. Up to four disk controllers and two communication line controllers also could be connected to the bus to communicate with the parallel processing units. The disk controllers supported mirrored disk pairs.

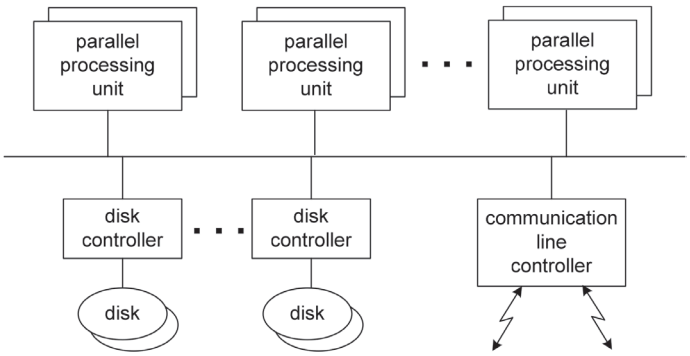

Parallel Computers

Each PPU was a duplexed fault-tolerant computer. A process in a PPU ran independently in each half. Periodically, the process pairs were synchronized. Thus, if one side of a PPU failed, the other side could carry on the processor functions in a manner transparent to the user.

Each PPU could contain one megabyte of memory.

In 1988, Parallel Computers was acquired by IMP, a British computer company. According to an IMP cofounder, IMP acquired Parallel Computers to gain entry into the U.S. computer marketplace.

Tolerant Systems

Tolerant Systems provided a series of System Building Blocks (SSBs) based on the National 32032 microprocessor. The SSBs could be arranged in a variety of ways by the user to achieve desired functions and levels of redundancy.

In a typical system, SSBs were used as communication servers to process requests from the users and to return responses, as application servers to process the transactions, and as file servers to access and update the database. In general, multiple servers of each type were available for redundancy.

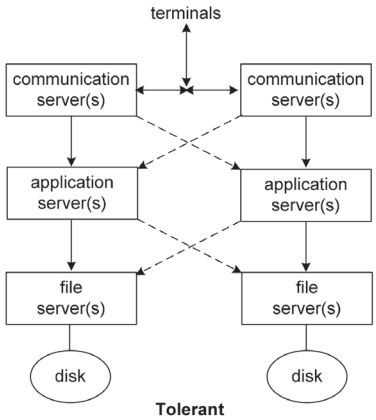

Tolerant

Fault tolerance was based on transactions. The communication servers maintained a log of all incomplete transactions, and the file servers maintained the before images of all incomplete updates. Should a server fail, its load was transferred to a like surviving server. All transactions being handled by the failed server were aborted, and the corresponding partial database updates were rolled back via the before images maintained by the file servers. Those transactions then were replayed from the communication server logs and enabled full recovery transparent to the user except for a time delay.

Tolerant Systems was renamed Veritas Software Corporation in 1989. After developing several new applications, Veritas went public in 1993 at a value of USD $64 million. Veritas subsequently was acquired by Symantec Corporation in 2005. Symantic produces software for security, storage, backup and availability

Summary

Following Tandem Computer's successful entry into the fault-tolerant server field, numerous companies attempted to follow. Of them, Stratus was the most successful. Others were absorbed or no longer exist.

After the success of Tandem and Stratus, IBM introduced its Parallel Sysplex fault-tolerant system in 1994. Today, most large mission-critical systems are powered either by Tandem (now HPE NonStop) or by IBM. Mission-critical systems on the edge (such as in sales offices) are largely managed by Stratus.

Much of the fault-tolerant technology we discuss here may seem woefully outdated. However, forty years after it was introduced, the Tandem NonStop architecture remains the same. That is simply amazing. Back when Tandem and similar systems were first introduced, there was no Internet. There was no Big Data. Cloud computing? Huh? From that perspective, what should we infer will be the scope of IT four decades from today? Will fault tolerance remain relevant or even exist? What we do know for certain is that the past is an indicator of future innovation and evolution. As such, appreciating the dawn of fault-tolerant computing will shape our anticipation of what lies ahead.

Dr. Bill Highleyman is the Managing Editor of The Availability Digest (www.availabilitydigest.com), a monthly, online publication and a resource of information on high- and continuous availability topics. His years of experience in the design and implementation of mission-critical systems have made him a popular seminar speaker and a sought-after technical writer. Dr. Highleyman is a past chairman of ITUG, the former HP NonStop Users' Group, the holder of numerous U.S. patents, the author of Performance Analysis of Transaction Processing Systems, and the co-author of the three-volume series, Breaking the Availability Barrier.

Back for More...

Richard Buckle >> **CEO** >> Pyalla Technologies, LLC.

Event season is in full swing and I am already planning my time away from the office so as to be able to attend. Bookended between GTUG and HPE Discover, there will be gatherings of the NonStop community in Texas (N2TUG), Southern California SCTUG) and a day of free education for BITUG users. While I am missing all the networking opportunities I will be attending both SCTUG and HPE Discover and I suspect as you read this column, some of these events will have already been held.

What is newsworthy for me is that the community continues to show enthusiasm whenever a major event is conducted. NonStop community members who subscribe to social media channels like Facebook and Twitter see the many updates from HPE, vendors and even users as they checked into flights to Berlin for GTUG.

While this next event may not appear on some calendars, I have to say I am impressed with how quickly Karen Copeland, along with members of the NonStop vendor community and Connect, turned a request to hold a partner summit (as raised at the 2015 NonStop Technical Boot Camp), into reality. As for me, I am looking forward to catching up with Karen, her team and the other vendors who will be participating. Gathering the vendors together at a time like this, when a number of us plan on heading to Las Vegas for 2016 HPE Discover, is very important and we owe our thanks to the HPE product management team.

Why am I as enthusiastic as I am about such an event? Looking back through the history of NonStop and indeed the user community, whenever ITUG held its major summit in San Jose, provision was always made to gather the NonStop vendor community together for a presentation "sampler" that more or less outlined what the then Tandem Computers product management would be covering. The intent was to make sure there wouldn't be any surprises during the summit itself and should other NonStop community members question vendors at their exhibition booths, they would be better informed and be able to more fully engage all those with questions.

As a former group manager of a number of product managers focused on communications and networking, there was always a demand from the vendors to know what was coming – remember, this was in the days prior to the overwhelming success of LANs, Client/Server computing and even TCP/IP. It may be hard to imagine but it was at one of these pre-ITUG updates where the decision was taken to move budget funds from SNAX to support this new TCP/IP networking and to check out potential support for the Internet.

And the decision wasn't taken lightly as the SNAX organization was a large group and consumed considerable funding. On the flip side of the argument however was the existence of a vendor prepared to spend their own money on providing a more modern SNA implementation on Tandem Computers and with its commitment, well, the rest as we like to say, is history. Tandem invested in TCP/IP and with it LANs and Client/Server computing and the enterprise customers needing SNA were able to turn to ACI Worldwide / InSession Inc. Fast forward to today and for the NonStop X systems, product management recommends uLinga from fomForte as its preferred solution for SNA support so in a way the tradition of engaging with vendors continues.

I am not predicting anything similar coming out from this 2016 edition of the partner summit but I am hopeful it will be a successful event. I am hopeful too that resources will be found to hold a separate day for vendors immediately prior to the upcoming 2016 NonStop Technical Boot Camp. After all, six months is a long time in the life of any technology or architecture and I am sure that there will be numerous interested parties to hear firsthand of progress being made with Hybrids, Yuma and yes, virtualization.

With solutions vendors already constructing cloud computing offerings based on the new NonStop X systems, and here I am writing about OmniPayments, Inc. and its OmniCloudX offering, whereby the OmniPayments' transaction-switch services are being offered from out of its cloud running on NonStop X. However, I suspect that even as OmniPayments may be first to market with NonStop X in the cloud, it may not be alone in this regard as x86 architecture lends itself to supporting such configurations.

While it's not HPE NonStop product management's place to comment on the likely future direction of this NonStop X family of systems, as a car guy, I can't help but draw comparisons with what market segments auto manufactures build to – BMW's 1 / 3 / 5 / 7 series (yes, ignoring the mutant even numbered participants) – calls for consideration that between the NonStop X 3 and 7 systems we may find a 5 at some point. Again purely speculation on my part, but if the 3 is based on a 1 and 2 core Xeon chip, and the 7 is optionally a 2, 4 and 6 core (perhaps more) Xeon chip, I could easily imagine there being further differentiation with the 5 coming with a choice of 2 and 4 cores and the 7 providing options for 4, 7 and more cores. Furthermore, just as BMW's 3 series is its best seller I expect the NonStop X 3 series to be HPE's best-selling NonStop X family member.

Truthfully? Bank this under the heading, mind you, of speculation but the point here is that for the first time in many years, the new NonStop X family comes with choice and with options. It will likely all be up to whatever budget is provided but with NonStop surpassing HP-UX as the Mission Critical Systems' Halo product, I am sure budget is being provided. And that too is why events remain so important as they remain the sole place where all in the NonStop community get to hear the official plan of record and where the likes of me can continue to dream up an even brighter future for NonStop. See you at 2016 HPE Discover!

SQLXPress

Not just another pretty face

An integrated SQL Database Manager for HPE NonStop.

Single solution providing database management, visual query planner, query advisor, SQL whiteboard, performance monitoring, MXCS management, execution plan management, data import and export, data browsing, and more...

With full support for both SQL/MP and SQL/MX

Intregrated with
XYGATE Merged Audit

Coming Soon!

Audit 100% SQL/MX & MP user activity

XYPRO
Mission Critical Security

Learn more at
xypro.com/SQLXPress

The guiding light for your mission critical business
Improve your NonStop'ness. Better always on!

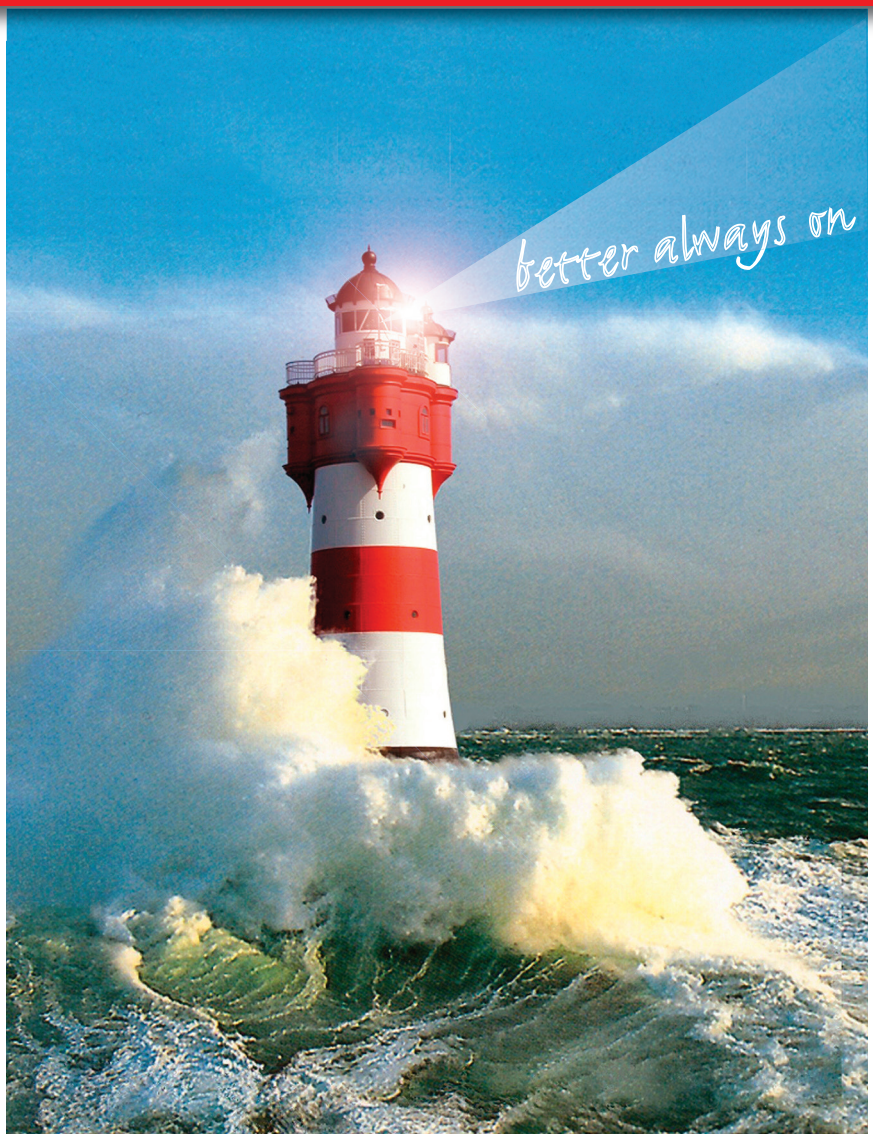

www.comforte.com/better_always_on

Today's demands of mission critical businesses and customers are ever increasing. Unreliable and unavailable systems and applications are not an option. Minimizing downtime whilst maximizing security and operational efficiency is therefore paramount for the IT department. If your light is going out, your business and your customers can get in trouble. Systems and applications can't stop; they must be on, always!

comForte „better always on" solutions help you gain …

Better Infrastructure
Make the most of best in class communications and connectivity solutions by providing end users and system administrators with high performance, secure and reliable access to NonStop systems.

Better Security
Protect your mission critical data-in transit and at-rest. Improve your overall security posture on NonStop and achieve compliance with industry standards and regulations.

Better Applications
Modernize your legacy applications from the database layer, through better integration in the enterprise all the way to refreshing the application's Graphical User Interface.

Better always on with comForte's unparalleled solutions for HP NonStop.

comforte®
better always on