

A BZ Media Publication

SD Times

● SOFTWARE DEVELOPMENT


YOU DREAM. WE DRIVE.

The world's data flows through our drivers.
Let us do it for you too. And free up your team to dream.

Progress.com/DataDirect



JANUARY 2016 • ISSUE NO. 321 • \$9.95 • www.sdtimes.com

 **PROGRESS** DataDirect



YOU CREATE. WE CONNECT.

Let us make the connections so you don't have to.
Oracle, MongoDB, Spark SQL. You name it. We do it all.

Get simple, standard SQL access to any
data source with Progress DataDirect.

Download your free trial today.

Progress.com/DataDirect

 **PROGRESS** DataDirect

SD Times

● SOFTWARE DEVELOPMENT

The Buyers Guide:
Application Performance Management



The Year of Containers

2015 saw Docker box up the market

Tips to get more
productive in 2016

Putting the auto
back in automobiles

Data entry errors. Consolidation headaches. Quality data shouldn't be a pain.

Try a little Vitamin Q.



We supply Vitamin Q – “quality” – in data quality management solutions that profile, clean, enrich, and match your customer data – and keep it healthy over time. Add Vitamin Q to your data integration, business intelligence, Big Data, and CRM initiatives to ensure accurate data for greater insight and business value.

**ACT
NOW**

Get a **FREE 30-day trial!**
www.MelissaData.com/vitaminQ



Solutions for
240+ Countries



10,000+ Customers
Worldwide



30+ Years
Strong



Data Quality &
Mailing Solutions



Cloud • On-Premise
• Services



Germany
www.MelissaData.de

United Kingdom
www.MelissaData.co.uk

India
www.MelissaData.in

Australia
www.MelissaData.com.au

MELISSA DATA®

Your Partner in Global Data Quality

www.MelissaData.com | 1-800-MELISSA

ASP.NET MVC!

The first cross-browser,
true WYSIWYG
HTML5-based
document editor.

Now with full
ASP.NET MVC support.



Give your users an MS Word compatible editor to create powerful documents and reporting templates anywhere - in any browser. Feature-complete including mail merge, sections, headers and footers, drawings and shapes, tables, barcodes and charts.

Available for ASP.NET Web Forms and MVC.

```
PM> Install-Package TXTextControl.Web
```

Live demos and 30-day trial version download at:
www.textcontrol.com



facebook



Microsoft
SharePoint 2010



Linked in



twitter

SEE THE WORLD AS A DATABASE

ADO.NET ▪ JDBC ▪ ODBC ▪ SQL SSIS ▪ ODATA
MYSQL ▪ EXCEL ▪ POWERSHELL



Visual Studio

Java

ODBC

SQL Server

Excel

BizTalk

MySQL

ODATA

Work With Relational Data, Not Complex APIs or Services

Whether you are a developer using ADO.NET, JDBC, OData, or MySQL, or a systems integrator working with SQL Server or Biztalk, or even an information worker familiar with ODBC or Excel - our products give you bi-directional access to live data through easy-to-use technologies that you are already familiar with. If you can connect to a database, then you will already know how to connect to Salesforce, SAP, SharePoint, Dynamics CRM, Google Apps, QuickBooks, and much more!



Give RSSBus a try today and see what we mean:

visit us online at www.rssbus.com to learn more or download a free trial.

rssbus

INTEGRATION YOUR WAY

Contents

ISSUE 321 • JANUARY 2016

NEWS

- 11 SD Times on the Web
- 14 News Watch
- 16 Forecast 2016: Predictions for the new year
- 18 Work alone? Get in the zone!
- 20 Keeping track of developer productivity
- 21 IBM adds cognitive APIs to its Watson platform
- 22 Diabetes app leads the way for open source in health
- 26 Android Studio 2.0 previews a shift in focus
- 26 NodeOS nears 1.0
- 28 HUD wants to put high-speed Web into low-income housing
- 31 TX Text Control X13 reaches out to ASP.NET MVC

COLUMNS

- 59 **CODE WATCH** by Larry O'Brien
Let deep neural networks help you
- 60 **GUEST VIEW** by John Chang
Looking behind the UI of mobile apps
- 61 **ANALYST VIEW** by Rob Enderle
2016: The year smartphones revolted
- 62 **INDUSTRY WATCH** by David Rubinstein
Seven habits of highly effective DevOps

FEATURES

Computer, take the wheel



page 40

Effective APM requires more data



page 47

2015: The Year in Review Containers take over the world



page 32

20 ways to build out your Azure deployment

page 54



Come to the Largest Wearable-IoT OEM Expo of 2016!

Learn how to design, build and develop apps
for the wearable technology revolution
at Wearables TechCon 2016!


**Wearables
TechCon**

July 18-20, 2016

San Jose Convention Center

The Heart of Silicon Valley

PLUS!

**Internet
of Things
TechCon**

www.WearablesTechCon.com





**Register
Early
and SAVE!**
**All-Access Pass
ONLY \$495**
Register before
Jan. 15, 2016

More than 60 Tech Sessions and Panels

- Product Design
- Software Design
- Leading SDKs
- Embedded System Development
- Application Development
- Electronic Engineering for Wearables Devices

...and more tricks and techniques that will set your wearable project apart!

Special Events

- Women in Wearables Luncheon
- Exhibit Hall Reception
- Industry Keynotes
- SDK Walkthroughs
- Lightning Talks
- Hackathons

"This is an excellent conference to attend to get an understanding from both a business and technical sense of Wearables and IOT."

—Kevin Jones, Electrical Engineer, Dept. of Defense

Who Should Attend

- Electronics OEMs
- Software Engineers
- Android Developers
- Design Engineers
- Software Developers
- Project Managers
- Hardware Designers
- Hardware Engineers
- Design Architects
- Software App Developers
- UI Designers
- Programmers
- Product Designers
- Embedded Developers
- Software Architects
- Product Managers
- System Architects
- Business Development Managers

Go under the hood of the hottest wearable devices!

A **BZ Media** Event





dtSearch®

Instantly Search Terabytes of Text

dtSearch's document filters support popular file types, emails with multilevel attachments, databases, web data

Highlights hits in all data types;
25+ search options

With APIs for .NET, Java and C++.
SDKs for multiple platforms.
(See site for articles on faceted search, SQL, MS Azure, etc.)

Visit dtSearch.com for

- hundreds of reviews and case studies
- fully-functional enterprise and developer evaluations

The Smart Choice for Text Retrieval®
since 1991

dtSearch.com 1-800-IT-FINDS

SDTimes

EDITORIAL

EDITOR-IN-CHIEF David Rubinstein
631-421-4158 x105 • drubinstein@bzmedia.com

SENIOR EDITOR Alex Handy • ahandy@bzmedia.com

**SOCIAL MEDIA AND
ONLINE EDITORS** Christina Mulligan • cmulligan@bzmedia.com
Madison Moore • mmoore@bzmedia.com

COPY EDITOR Adam LoBelia • alobelias@bzmedia.com

SENIOR ART DIRECTOR Mara Leonardi • mleonardi@bzmedia.com

COLUMNIST Larry O'Brien

CONTRIBUTING WRITERS Alyson Behr, Rich Crane, Patrick Hynds,
G. Arnold Koch, Lisa Morgan,
Alexandra Weber Morales

CONTRIBUTING ANALYSTS Rob Enderle, Michael Facemire, Mike Gilpin,
Mike Gualtieri, Jeffrey Hammond, Al Hilwa

CUSTOMER SERVICE

SUBSCRIPTIONS subscriptions@bzmedia.com

ADVERTISING TRAFFIC Mara Leonardi
631-421-4158 x109 • mleonardi@bzmedia.com

LIST SERVICES Shauna Koehler
631-421-4158 x112 • skoehler@bzmedia.com

REPRINTS Stacy Burris
631-421-4158 x108 • sburris@bzmedia.com

ACCOUNTING Viena Ludewig
631-421-4158 x110 • vludewig@bzmedia.com

OTHER QUESTIONS 631-421-4158 • info@bzmedia.com

ADVERTISING SALES

PUBLISHER David Lyman
978-465-2351 • dlyman@bzmedia.com

WESTERN U.S., WESTERN CANADA, EASTERN ASIA, AUSTRALIA, INDIA
Paula F. Miller
925-831-3803 • pmiller@bzmedia.com

EASTERN U.S., EASTERN CANADA, EUROPE, MIDDLE EAST, WESTERN ASIA
Jonathan Sawyer
603-924-4489 • jsawyer@bzmedia.com

PUBLISHING DIRECTOR Ted Bahr
631-421-4158 x101 • ted@bzmedia.com

BZ Media

PRESIDENT Ted Bahr

FOUNDING EDITOR Alan Zeichick

BZ MEDIA LLC

225 Broadhollow Road, Suite 211
Melville, NY 11747
TEL 631-421-4158
FAX 631-421-4130
www.bzmedia.com
info@bzmedia.com

Realize your Vision

Award winning modeling & design tools for Software, Business & Systems



Model Design Simulate Build Manage



ALM and
Development
Tools



Collaborative modeling & design environment
Analyze, design & build with exceptional quality & efficiency
UML[®] based core, plus many standards based extensions
BPMN[™], SysML, SoaML[™], BPEL[™] & many others
Simulation, coding, debugging & visualization tools
Performance, price & precision
Full lifecycle support with end-to-end traceability

Join the community www.community.sparxsystems.com

SPARX
SYSTEMS

www.sparxsystems.com

UML | BPMN | DDS[™] | BPEL | SysML | TOGAF[®] | C++ | .NET | Java[™] | ASP | PHP | Delphi | SoaML[™] | ArcGIS[™] | SOMF[™] | XML[™] | XSD | XMI[®] | WSDL | *More

DocuVieware

NEW Universal HTML5 Viewer & Document Management Kit



zero-footprint solution



super-easy integration



fast & crystal clear rendering



fully customizable UI
look & feel



mobile devices optimization



annotations, thumbnails
bookmarks & text search

supports nearly 100 formats



George Zimmer dons some new threads

Remember the Men's Warehouse guy who was shockingly ousted from his own company? He's back, and he's taking some tips from Valley startups like Uber with his own startups: zTails and Generation Tux. "While Zimmer's old business had traditional enterprise IT, [executive vice president of zTails and Generation Tux Matt Howland] said that both zTails and Generation Tux are 100% cloud-hosted. 'We use Salesforce pretty aggressively on the zTails side. We leverage Heroku, but on Generation Tux we do Docker and whatnot,' he said." So far, George Zimmer likes the way his companies looks, and you can see why at bit.ly/1NgMFeV.

My first computer

A bit of a nostalgia trip on the Web happened when Computerphile asked computer scientists what their first computer was. The Dragon 32, the Armstrad 286 and the Commodore PET were brought up, among many others (including a simple mainframe). Their recollections were put in a video, available at bit.ly/1QuOUzx. (For the record, mine was an IBM PC compatible.)



Safe from without, not within

Usually you hear about hacks and intrusions that occur from the outside. In turn, that usually means you're not hearing about the vulnerabilities—malicious or accidental—from those on the inside. Wynyard Group's Derek Brown has something to say on this: "According to a 2014 CERT study, 28% of all security incidents occur as a result of an insider—either from an employee or from an employee who allows an external threat into the network. The question, therefore, is how can organizations best defend themselves against both external and internal threats that are evolving and have the potential to cause significant reputational, financial and operational damage?" He asks and answers this at bit.ly/1TGb2Ew.



The early test catches the defect

Gaps in software are to be expected, and developers are bound to miss something when they're coding. Testing can catch these issues, but wouldn't it be better to catch them *before* they become a nuisance? KMS Technology's Mush Honda thinks so: "By including testers from day one, you reduce the chance of misinterpretation and give yourself an opportunity to spot potential gaps or misunderstandings. But the communication and assessment has to be continuous throughout the project; the second it slips, you're running the risk of a gap opening up." If this sounds like something your business can get into, you can read more at bit.ly/1SXeyVM.



SD Times wants to hear from you.
Join us on **LinkedIn** and **Facebook**.



Open Create Convert Print Save
files from your applications!



Contact Us:

US: +1 888 277 6734

EU: +44 141 416 1112

AU: +61 2 8003 5926

sales@aspose.com

SCAN FOR
20% SAVINGS!



BUSINESS FILE FORMATS



ASPOSE.Cells

XLS, CSV, PDF, SVG, HTML, PNG
BMP, XPS, JPG, SpreadsheetML...

ASPOSE.Words

DOC, RTF, PDF, HTML, PNG
ePUB, XML, XPS, JPG...

ASPOSE.Pdf

PDF, XML, XSL-FO, HTML, BMP
JPG, PNG, ePUB...

ASPOSE.Slides

PPT, POT, POTX, XPS, HTML
PNG, PDF...

ASPOSE.BarCode

JPG, PNG, BMP, GIF, TIF, WMF
ICON...

ASPOSE.Tasks

XML, MPP, SVG, PDF, TIFF
PNG...

ASPOSE.Email

MSG, EML, PST, EMLX
OST, OFT...

ASPOSE.Imaging

PDF, BMP, JPG, GIF, TIFF
PNG...

+ MANY MORE!

Get your FREE evaluation copy at www.aspose.com

.NET

Java

Cloud

NEWS WATCH

Netflix engineers show how to do Linux performance analysis

The Netflix Performance Engineering team has written a guide for finding performance issues within 60,000 milliseconds.

Netflix has a EC2 Linux cloud, which is constantly monitored by performance-analysis tools to make sure everything is running smoothly. These tools include Atlas (for cloud monitoring) and Vector (for on-demand instance analysis).

The Netflix Performance Engineering team wrote that these tools help them solve most of their issues, but sometimes they “need to log in to an instance and run some standard Linux performance tools.”

By following 10 commands, the team said in 60 seconds an engineer can get an idea of system resource usage and running processes.

A detailed list of the commands is available on Netflix's blog at nflx.it/1HC6q1C.

Apache Kafka updates to 0.9

Apache released version 0.9 of Kafka, adding support for SSL for the first time, as well as new consumer APIs.

This latest version of Kafka includes updates that should make it easier to integrate this high-throughput distributed messaging systems into data-driven architectures. This release is the first to include the newly rewritten consumer client, which allows developers to use the deeper features of Kafka without having to use Java.

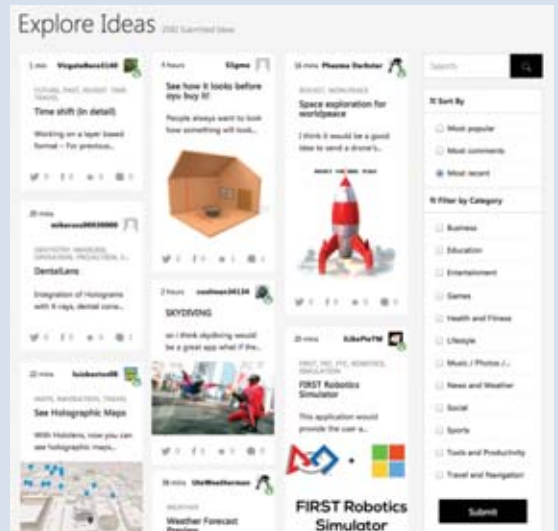
The Kafka consumer client is also much thinner now,

Microsoft wants your HoloLens ideas

Microsoft wants to start bringing holographic computing experiences to consumers, but it needs some help getting there. The company has announced the “Share Your Idea” campaign to learn what developers are working on and to jumpstart their ideas.

Through the campaign's website, developers can share what experiences they'd like to see on the HoloLens, and the HoloLens community will be able to vote on the best ideas. Microsoft will pick three of the most popular ideas, then host a final vote on Twitter. The HoloLens team will then help make the winning idea a reality, open-sourcing the code so other developers can try it and build on it.

According to Microsoft, when thinking of an idea, developers should consider if the hologram can stay in place, scale, interact with users, utilize sound, and work with surfaces in the user's environment. Developers have until Jan. 11 to submit ideas.



allowing for fewer dependencies and less overhead. Another change to Kafka's consumer client API allows multiple Kafka nodes to better run as a group thanks to the addition of a highly available coordinator node to the herd.

Adobe moves away from the Flash brand

Adobe Flash is soon to be no more. The technology, which once was ubiquitous on desktops and browsers, will now be called Adobe Animate CC (for Creative Cloud) to reflect the changing times.

The change comes as the Web moves more toward open Web and HTML5 standards. With the emergence of HTML5, the need for plug-ins like Adobe Flash ended and gave companies the ability to add rich features into browsers without having to install software beforehand. Big names such as Mozilla, YouTube and Amazon have already started to move

away from the technology.

The company is expected to move to Adobe Animate CC this year. The new product will continue to support Flash and AIR formats, and will continue to improve its HTML5 capabilities over time.

In addition to renaming the product, Adobe is working on new features. Those include improved audio syncing, enhanced color changing, a 360-degree rotatable canvas, vector art brushes, Creative-Sync integration, multiplatform support, 4K+ video export, and OAM support.

Progress announces major overhaul of the Telerik Platform

Progress recently announced Telerik Platform 2.0, its mobile app development platform for Android, iOS and Windows Phone apps. The latest release features a major overhaul to the platform's look, feel and UX.

According to Brandon Satrom, general manager for developer platforms and tools at Progress, the update provides more than just a new design; it provides a unified UX focused on helping developers simplify their workflow by making it easier to leverage critical services.

The release also aims to help enterprises connect their mobile apps to existing legacy data. Developers can leverage data in MySQL, Oracle, PostgreSQL, Salesforce and SQL Server using a series of connectors powered by Progress DataDirect drivers.

Other features include the beta release of CI build support with a Jenkins plug-in, the latest updates to NativeScript, a Verified Plugins Marketplace for NativeScript, enhancements to Telerik Backend Services, and support for Android 6, Apple Watch and iOS 9 apps.

Facebook is open-sourcing its AI hardware

Facebook has announced plans to open-source its artificial intelligence hardware design. The Open Rack-compatible hardware, codenamed Big Sur, is designed to handle AI computing at a large scale, according to the company.

Facebook hopes open-sourcing the hardware will make it easier for AI researchers to share their techniques and technologies, and improve on the system. In addition to open-sourcing Big Sur, Facebook plans to submit its design materials to the Open Compute Project, a collaborative initiative to deliver "efficient server, storage and datacenter hardware designs for scalable computing," according to the project's mission statement.

Big Sur leverages NVIDIA's Tesla Accelerated Computing Platform, and features eight high-performance GPUs, improved performance from the company's previous generation, and optimized servers for thermal and power efficiency.

Sencha GXT 4 covers Java developers building Web apps

For developers who use tools to build desktop applications, there's Sencha's announcement of the availability of Sencha GXT 4, which now allows them to extend their desktop applications to tablet devices with little to no modification to the existing application codebase.

Sencha GXT 4 is a comprehensive Java framework for building feature-rich applica-

tions, which Java developers can use to build and deploy HTML5 applications for desktops and tablets, as well as convert existing GXT applications to mobile. The framework also includes support for RPC, RequestFactory, and JSON so developers can feed data to the applications using any data source.

Cloud Foundry unites worldwide

Cloud Foundry has developed a certification program designed to keep the core of this Platform-as-a-Service project compatible within cloud providers. No matter the infrastructure below, said Cloud Foundry's CEO Sam Ramji, "The core, what we own as a Foundation, is all interoperable from cloud to cloud."

This certification ensures that the infrastructure below a hosted version of Cloud Foundry is "app compatible," said Ramji. "We have demanded that every company saying they're using Cloud Foundry commercially has to use the upstream [code]."

The cost to create a product labeled "Cloud Foundry" is US\$50,000 per year, but members of the Foundation are offered a discount.

Gradle raises \$4.2 million

Gradle Inc., the company behind the open-source enterprise build automation system, has announced a US\$4.2 million round of funding to expand its company and system.

According to CMO Miko Matsumura, the company has seen 30x and even 100x improvements in customer and partner

release cadence and build times. Going forward, the company plans to add more parallelization, distributed caching, a new model, and deeper IDE integration with Gradle.

The company will also work on its commercial SaaS offering, Gradle.com, to help enhance developer productivity. Gradle Inc. will work on new infrastructure and tools to help developers build and deliver complex software in a faster and more cost-effective way.

Google releases beta of AngularJS 2.0

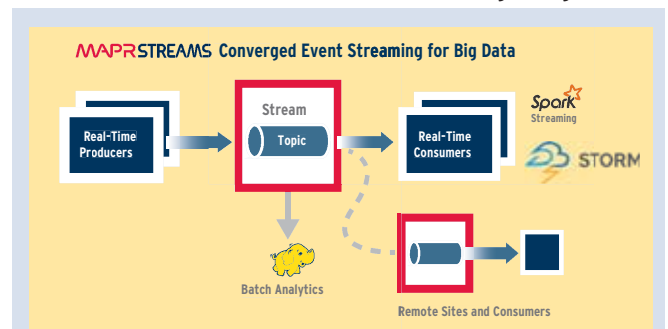
AngularJS fans will have to get used to a lot of new features and rework their applications for the forthcoming release of Angular 2.0. This

release was deemed necessary to unleash the JavaScript application framework from architectural limitations.

Some changes between version one and two can be intermingled thanks to a free tool from the Angular team inside Google. This will allow the use of Angular 2 idioms inside Angular 1 applications.

Other changes, however, require the entire application to be reworked in Angular 2. These stem from the decision to decouple Angular from the DOM.

More features are coming as the development process continues. A command-line interface is still in the alpha stages, and the binary size of Angular overall is being trimmed to allow mobile usage to be quicker and more lightweight. ■



MapR introduces publish/subscribe streaming for Hadoop

While many in the Big Data space are talking about stream processing, MapR announced the availability of Streams, a new product in its Hadoop stack that can be used to stream events across clusters distributed around the world. Streams offers a publish-and-subscribe model for event-driven data access and decision-making.

Streams can handle the event-driven needs of an entire data architecture. That means it can run those event checks against data in Spark, or prepare it for a system like Apex or Kafka. The company said this can help bring a single view to all of those separate methods of data analysis by hooking them into the worldwide event stream across multiple Hadoop clusters.

Streams is currently in limited preview with select clients, but will be generally available to all customers as a part of the MapR platform. It will also be a part of the Community Edition of MapR.



To paraphrase that great thinker, Ferris Bueller: "Technology moves pretty fast. If you don't look around once in a while, you could miss it." So, to get 2016 rolling, we've asked a number of luminaries and thought leaders in the software development space to look around and tell us what they expect from the field this year.

Kelly Stirman, VP of strategy

MongoDB

Kafka will become an essential integration point in enterprise data infrastructure, facilitating the creation

of intelligent, distributed systems. With the growth of IoT, global deployments and microservices, the need to capture and control in-flight data before it's stored in a database is becoming more important. Kafka and other streaming systems like Spark and Storm will complement databases as critical pieces of the enterprise stack for managing data across applications and data centers.

Grady Booch, Fellow

IBM

There exists a temporary plateau in the advance of software development, for there are many things

we know how to do well: Continuous Integration; doing the simplest, testable thing possible; the role of stable interfaces; the place of new languages such as Swift as well as the maturation of older languages such as C++ and Java, to name a few. However, despite our best practices, the very nature of the systems we must build are changing in

subtle ways, and that will propel the art and the science of software development to the next level. Security, of course, will always remain an issue, and as we integrate billions upon billions of devices into our world, security must take a more primary place in development. The biggest unknown, and therefore the biggest opportunity for innovation in software development, is the development of systems that learn. Herein we face challenges in how best to shape the life cycle of systems whose behavior is shaped not just by software but by the structured and unstructured knowledge they ingest. Rather than just cut code, we must now find the best practices for teaching our systems. I suspect we'll see a new class of jobs emerge: not just machine-learning professionals, but perhaps even software therapists.

Monte Zweben, cofounder and CEO

Splice Machine

Spark will kill Map Reduce, but save Hadoop. Map Reduce is quite esoteric. Its slow, batch nature and

high level of complexity can make it unattractive for many enterprises. Spark, because of its speed, is much more natural, mathematical and convenient for programmers. Spark will reinvigorate Hadoop, and in 2016, nine

out of every 10 projects on Hadoop will be Spark-related projects.

Andrew Levy, founder and CEO

Crittercism

Next year, both Marriott and Hilton will have wide rollouts of keyless entry for rooms you unlock your

door with your smartphone. We'll eventually see other sectors follow suit, such as car rentals, that will continue to fuel growth for the e-commerce category. As such, app developers should be prepared to find ways to automate and streamline other daily functions with the use of a smartphone.

Erik Sebesta, Chief Architect and Technology Officer

Cloud Technology Partners

The major cloud providers will grow at 100+%. The rest of the world will wish they were the major cloud providers.

Rajiv Gupta, cofounder and CEO

Skyhigh

Companies will start to pay off cloud security debt. More and more companies are full-speed ahead on cloud, but so far security has lagged behind. There's a gap between where cloud security budgets currently are

and where they should be based on overall security spending. According to Gartner, companies allocate just 3.8% of cloud spending to security, compared to 11% from overall IT budgets. In 2016, budgets for cloud security will outpace overall IT security spending as companies play catch-up.

Bill Curtis, executive director



Consortium for IT Software Quality

While software glitches will drive up the demand to identify and reduce IT risk, the

problem will only get worse until the C-Suite understands the financial trade-offs of business IT taking into account software quality measurement. Last year's headlines show that businesses do not understand the level of risk in their software. That's because they lack industry benchmarks and best practices for evaluating the quality of their software. It's often the structural quality of the software—the non-functional characteristics such as reliability and security—that hackers exploit.

Richard Soley, executive director



Industrial Internet Consortium

We're going to see the Internet of Things go into the trough of disillusionment, but I

suspect it's going to be a really shallow trough, because there are already design wins/case studies to talk about. At the recent IoT Solutions World Congress, business leaders, operational leaders and CIOs gathered to listen to more than 120 speakers in 83 sessions discuss the opportunities and challenges in deploying the Industrial Internet of Things. This conference...featured global leaders who are deploying solutions today. Among them are a smart building that Schneider Electric collaborated on; wearable technology that Massachusetts General Hospital is using to monitor research patients' stress levels; or a car that was designed via crowdsourcing and was created using 3D printing in less than 24 hours.

Mathew Lodge, coo

Weaveworks

Container Management starts to look like a "Game of Thrones" wedding scene. In "Game of Thrones," blood runs as freely as wine at weddings, where feuding families all seek control of key strategic territory. For production Docker deployments, the strategic territory is management and control of container fleets, and it's getting increasingly crowded and contentious. Google was early with Kubernetes, based on its own experiences managing a billion container-starts per week. Apache Mesos is a popular alternative in use at large-scale players like Apple, and you also have Marathon from Mesosphere. AWS introduced its own container orchestrator at re:Invent for its EC2 Container Service. That's too many alternatives, too soon: 2016 is going to be brutal for vendors, and potentially confusing for users.

Cole Crawford, cofounder and CEO

Vapor IO

Open-source code will replace open interfaces. With the acceleration of permissive licenses and their adoption rate, open washing will result in open shamming. Upstream code contributions will be increasingly managed via top-down governance structures and won't be limited to technical merit and usability. Open-source code will also continue to replace open interfaces as companies learn that the latter still allows for vendor lock-in and gratuitous differentiation.

Neelan Choksi, president and COO



Tasktop

In 2016, it will no longer be the sole purview of CIOs and CTOs to have insight into software operations, but also the CEO, CMO, COO and CFO. Anyone in those roles who doesn't know what is going on with their software will be at risk of losing their jobs. So expect to see more CEO firings in 2016 because of software failures. On the flip side, there will be more tech/software-savvy CEOs coming on the scene driven by software suc-

cesses. The bottom line: The C team will be expected to have a better understanding of how software drives the company's business so that they can help make the right decisions.

Hung Nguyen, founder and CEO



LogiGear

We've gone from DLL to COM objects to custom controls to Web objects and now micro-services. We're not sure you how testing will make sense. It's more about how the process will work, who will do what kind of testing and when.

Fima Katz, cofounder and CEO

Appery.io

Enterprise mobility will evolve from ad hoc projects to broader mobility and digitalization projects. Enterprises that built their first mobile apps in 2015 will look to create more apps across all departments in order to simplify and streamline business. From sales and marketing to customer services and manufacturing, increased mobilization will be key in making business easier and more efficient. Further, in 2016, we'll begin to see enterprises evolve their mobility efforts into plans for broader digitalization. Enterprises in key verticals such as healthcare, payments, automotive and energy will start to experiment with IoT devices. These efforts are going to require adoption of a more complete mobile integration platform.

Steve Goldsmith, general manager of HipChat

Atlassian

2016 will be the year that ChatOps expands from software teams to all business teams. Software is eating the world, and teams outside of software have seen this and are taking some of the best techniques used by software teams like Uber, Netflix and Twitter into the non-technical world. ChatOps is the operating system for great teams. It is the place where the right work, people and tools come together. In 2016, marketing teams, IT teams, legal teams and HR will start to adopt this model of work. ■



Work alone? Get in

Productivity tips to get more done on your

BY ALEXANDRA WEBER MORALES

In today's "gig economy," a growing number of workers are freelancers, entrepreneurs or remote employees. Being your own boss isn't always intuitive, however, so we asked solo developers, startup founders and other tech types for tips on how to get in the zone when you work alone.

In the early days of telecommuting, the commute-free home office was many a "solopreneur's" dream. But just as paper has taken second place to reading on the Internet, joining a coworking space is increasingly seen as a productivity imperative.

"I've tried working in a number of different places—coworking space, home office, and coffee shops—and I've found that coworking space is the most productive for me," said Lingke Wang, CTO and cofounder of a startup called Ovid. "The reason is two-fold. First, I prefer dual monitors plus an external keyboard and mouse when programming. I like to have one screen where I code and another where I can view the application, run tests and read documentation. This alone increases my productivity two-fold beyond my one-screen laptop. Furthermore, when I type for an extended amount of time on my laptop keyboard, my wrists will begin to hurt. My ergonomic external keyboard helps solve this problem."

In addition to having a custom setup at his coworking space, Wang appreciates the focus and higher energy than he had in his home office. "I used to be lazier because home is where I relax. I eventually moved to a coworking space where I can fully separate work and personal life," he said.

John Lin, head of product for another startup, Roomi, notes that "A coworking space can cost upwards of US\$400 a month per desk in popular startup cities with expensive rents, such as San Francisco and New York. But the most valuable intangible of a coworking environment is the community that you're surrounded by on a daily basis. Everyone in these spaces is going through the struggle together, which can be both motivating and inspiring."

Keep unusual hours at home

For some, the home court advantage is real. David Mercer, founder of SME Pals, says he works quickest alone. "I also prefer to work at strange times. If I have difficult problems to solve, I like to pace around or go for walks, which tends to be frowned on in formal corporate environments," he said.

A novel idea is to use your living space as a coworking hub for others. "Even though they work for different companies, I invite friends over to work from

my house. It creates a social environment and keeps you accountable," said Charlie Cohn, head of marketing at StudySoup.

The remote-work setup is essentially an engineering problem, so treat it like one. "For me, being at home in a well-structured environment works best, but my colleague who does back end prefers to be in a busy coffee shop. Just be lean, try, test and iterate until you find your own best solution," said Vitor Avelino, tech lead and cofounder of Bounty Most Wanted.

As anyone who has read "The 4-Hour Workweek," by Tim Ferriss, or "Getting Things Done," by David Allen, knows, e-mail is a major barrier to productivity and a source of infinite interruptions. To get in the zone, Rogelio Triviño, CEO of Shelfic, disconnects e-mail and incoming messages from social media. "Check messages once or twice a day, using these moments to disconnect from your tasks and with a time limit. Order incoming requests on a notes script. I use Google Keep for tasks, lists and to-dos," he said.

Sean Tepper, author of "Earn More, Work Less, Live Free," agrees that checking e-mail twice a day in 30 minutes or less per session is the ticket. "By dedicating specific time slots for checking e-mails, you will actually be able to cruise through more e-mails in less time," he said. "By removing e-mail-



the zone!

own in 2016

checking from the core part of your day, you will be able to dedicate your time to higher priority, uninterrupted, billable tasks. It will be tempting to open your inbox, but I strongly advise not to. Leave it aside. If a client has an emergency, that is what a phone is for.”

Jeff Gambera is an Oregon-based DevOps and quality hacker at the open-source machine-learning company H2O.ai. “Working remotely requires a different approach to pacing and self organization,” he said. “Remember that you can’t rely on the visual and social cues of when people go for lunch or other socialization, so set aside a few moments for yourself.”

He also recommends getting a robot or other telepresence device so you can virtually “go to where the action is.” But you might want to wait until virtual reality improves.

“We proposed the iPad robots but didn’t do them,” said Greg Beckman, director of engineering for OpenText Core, an enterprise information management SaaS out of Waterloo, Ont. “I think that telepresence—the tools for that unfortunately kind of suck. I could Skype you or we could FaceTime, but if we have a meeting where people are jumping up and standing around a whiteboard—it’s hard to have that communication to someone sitting in front of a monitor.”

Don’t fear the phone

Beckman manages a team of 60 across Canada and the U.S. Of his 15 to 20 remote technical workers, only about five work from home, he said. He finds it critical to bring remote workers into nearby offices as often as possible, and to keep communication channels open. But he has found that millennial employees can be reluctant to use the channel invented by Alexander Graham Bell.

“I’ve got a very young team, not as many senior developers, and with that can come a reluctance to even just pick up the phone and use other forms of communication rather than just chat and IM. We’ve had a couple situations where the personal touch has lacked,” said Beckman.

Though such agile practices as standup meetings and two-week iterations have become commonplace, pair programming is less common. For some remote workers, it’s a life-saver nonetheless, said Steve Gibson, director at the Web company JotForm.

“In theory it means you’ll get half as much work done,” he said. “Advocates of it say that the code ends up more robust, is produced faster, and both developers learn from each other as they go. For those working outside an office, it also gives much needed social interaction.”

If you’re a fan of keying your work

activities to your energy levels and moods, one suggestion is to embrace a form of multi-tasking that pits your avoidance behaviors against each other: “A recipe to not procrastinate is to run several projects at the same time, going every day with the one that fits your mood,” said Triviño, who is working toward the launch of shelfic.com and writing a software development buyers guide.

Become a slacker

Google Docs or DropBox for cloud-based document management, JIRA for issue tracking, Skype and FaceTime for messaging, virtual assistant apps and agencies—these are just a few of the plethora of productivity tools now available to remote workers. But one consistently came up among the developers interviewed: Slack. At the most basic level a simple chat tool, Slack also has automation channels to integrate notifications from various workflow systems.

“Slack is one of those Silicon Valley darlings right now. We’ve used other tools like Microsoft Link or a home-grown one, but we’ve found that Slack is more seamless. It just ties everything together,” said Beckman, who notes that the tool can also help cement group culture. “Certain commands pop up a meme; the Slack channel will Google search an image for ‘ship it.’”

Andrey Khusid, CEO of RealtimeBoard, a virtual whiteboard service, concurs: “We have several corporate Slack channels separated by topics, and each employee has installed it in his smartphone, so notifications come in real time. If you are a freelancer, you can join the Slack group of the most relevant community.”

However you earn your bread in the gig economy, it’s critical to find ways to interface with other carbon-based life-forms—here in the third dimension, not virtually. “It is extremely valuable to share ideas with others,” said SME Pals’ Mercer. “Often people will bring new perspectives, mention new technologies, apps, trends or articles that can really help you. Don’t stay isolated for long periods of time, and online forums, Google, and YouTube don’t count.” ■

Keeping track of productivity

What's moving developers forward and what's holding them back?

BY GAIL MURPHY

Talk of software development productivity abounds. New languages, like Dart, promise software developers that they don't have to choose between productivity and performance for the programs they write. New DevOps tools, like ElectricFlow, promise to accelerate everything to improve developer productivity. Perhaps your team's productivity is not limited by languages or tools, but rather how you organize or track your team's work, communication, or something completely different.

Do you know what most affects software development productivity in your organization? Probably not, and you are not alone.

Along with my colleagues André Meyer and Thomas Fritz at the University of Zurich, and Thomas Zimmermann at Microsoft Research, we searched for what is known about software development productivity. We found many theories about how productivity might be measured and what organizational factors might affect productivity. We also found many suggestions for how to improve productivity. But we found nothing about how software developers themselves think about productivity and what they perceive affects it.

What's going on?

To gain a broad view about what productivity means to software developers and how developers think about assessing their own productivity, we conducted a survey of 379 developers whose average professional experience was 9.2 (± 7.3) years.

Not surprisingly, the two top factors leading to a productive day were when developers complete tasks (53% of respondents) and when they are able to work with few interruptions or distractions

(50%). Also not surprising was that developers are productive when they have no meetings (21%). The fourth and fifth most mentioned factors indicated the reflective nature of many developers' work: Developers perceived they were productive when clear goals are set (20%) and they plan their workday (17%).

At a more fine-grained level, we asked developers which activities they



perform are productive or not. They reported coding as the most productive activity (71%), and most (58%) thought meetings were unproductive, particularly when they lacked goals or had too many attendees. But almost a fifth (17%) thought meetings with clear decision-making were productive.

Developers were in general satisfied with their productivity on the previous workday and workweek. Most (78%) assessed their productivity based on tasks completed, although several (27%) mentioned such measures as lines of code, number of commits, number of bugs found or fixed, and e-mails sent.

To further deepen our understanding of developers and their productivity, we observed 11 developers from three companies of varying sizes for four hours each on a single workday. The observer recorded everything the developer did: tasks performed, tools used, time spent using each tool and so on. We then analyzed the data to understand what work

occurred, when, and for how long.

In these four-hour sessions, developers worked on two to 10 tasks. Work on tasks was not consecutive; in fact, developers switched frequently between tasks with a mean switch rate of 13.3 (± 8.5) times per hour. The average time spent on a task was 6.2 (± 3.3) minutes. Despite all these switches, developers mostly (73%) reported the sessions as productive. These task switches happened for many reasons, including helping coworkers to unblock themselves, as well as because they themselves were blocked—for instance, waiting for a build to complete.

While performing tasks, developers were involved in many activities. They switched activities very frequently: Forty-seven (± 19.8) times per hour, spending an average of 1.6 (± 0.8) minutes on each activity before switching. The largest amount of time was spent on coding (32%), followed by testing (12%).

What can one do with this information? Building on quantified self-movement, we have started to work on dashboards that bring together keyboard activity and activity in software repositories with developer input about when they are productive. With such information, we can help developers reflect on whether they are most productive in the mornings or afternoons, and we can start to build tools that help them protect that time by automatically adding calendar entries to block their time as well as other simple aids. We are also working on methods to assess the impact of task switches that are disruptive to try to help developers organize their workday to stay focused on tasks requiring deep thought or creativity.

There likely is no one magic tool or technique that one can wave to suddenly improve developer productivity. But by understanding what your developers do all day and what activities they find full of friction, you can take steps to keep developers in the flow and producing great software. ■



Gail Murphy is chief scientist and cofounder of Tasktop.

FREE COMMUNITY LICENSE

A \$9,975 VALUE FOR FREE | SUPPORT AND UPDATES INCLUDED

CLAIM YOUR LICENSE

syncfusion.com/SDTimescommunity

Comprehensive offering includes more than 650 components across 13 platforms, an easy-to-use big data platform, and much more.



WEB

ASP.NET
ASP.NET MVC
HTML5/
JavaScript
Silverlight
LightSwitch



MOBILE

Android
HTML5/
JavaScript
iOS
Universal
Windows
Platform
Windows Phone
WinRT
Xamarin



DESKTOP

Windows Forms
WPF
Universal
Windows
Platform



FILE FORMATS

Excel
Word
PDF
PowerPoint



DATA SCIENCE

Big Data Platform
Predictive Analytics



ENTERPRISE SOLUTIONS

Dashboard
Platform
Report Server

IBM delivers cognitive APIs for IoT through Watson computing platform

Natural-language processing, machine learning, analytics bring context to data

BY DAVID RUBINSTEIN

IBM announced it is making cognitive Web services a part of its Watson IoT Foundation with the addition of four API families to the platform.

The API families deliver natural-language processing, machine-learning capabilities, text analytics, and video and image analytics, according to the company's announcement. Those APIs are among the 28 IBM already publishes for Watson computing, and the company is making the four directly referenceable so Internet of Things systems can combine them to gain context surrounding the information, Chris O'Connor, general manager of IoT at IBM, told SD Times.

"IoT starts with the device, but it's also about the environment around it," he said. Offering the example of an Airbus airframe, he said there are 6 million parts in the airframe provided by 3,000 suppliers. (An airframe is the mechanical structure of the plane, minus the engine.) Each of those parts could be an instrumented IoT instance, he said, that fit together and relate to each other. But environmental factors affect the plane in the air, such as weather, wind speed and turbulence, and these external variables drive demands on those IoT devices. So, he said, "You need to bring more data together to gain context, and Watson can make the data contextually relevant" using the capabilities underlying the APIs.

According to IBM, the Natural Language Processing APIs enable digital solutions to understand human language and its intent. The Machine Learning Watson APIs perform data processing, monitoring data inputs, and how the data is being utilized to create rankings of the data based on priorities it has learned. The Video and Image Analytics APIs enable monitoring of



From left: David Kenny, CEO, The Weather Company; John Kelly, SVP, solutions portfolio and research, IBM; Harriet Green, GM, IBM IoT & Education; Laurent Martinez, SVP, head of services business unit, Airbus; and Matthias Rebellius, CEO, Siemens Building Technologies, at the opening of IBM's Watson Internet of Things global headquarters in Munich, Germany, on Dec. 15.

unstructured data from videos, photos and other images to identify patterns.

Here's the example IBM provided: Video analytics that monitor security cameras might detect a forklift in a restricted area, creating a minor alert in the system. Three days later, an asset in that area begins to exhibit decreased performance. The two incidents can be correlated to identify a collision between the forklift and an asset that might not have been readily apparent from the video or the data from the machine.

Finally, the Text Analytics APIs enable data mining for unstructured information such as call center calls, maintenance logs, blogs and tweets to find correlations and patterns. So, phrases in tweets such as "The brake pedal feels mushy" can be correlated to other entries to determine if there are issues in a particular make and model of a car.

IoT, O'Connor said, is advancing from the simple collection of data from

myriad devices to more complex system processing, which can enable business optimization, the creation of new business models, and businesses to connect directly to customers who buy their products. In the example of automobiles, manufacturers and drivers have a layer between them: dealers. Advanced cognitive systems can help manufacturers communicate directly with buyers.

The Watson IoT Foundation is cloud-based with datacenter presence points in 40 centers around the world, O'Connor said, enabling connectivity, information management, security and cognition "just about everywhere." Further, IBM announced that it is opening a global headquarters for the Watson Internet of Things unit in Munich, Germany, which will also serve as a research center and innovation lab.

"Quarter to quarter, we're seeing a spike in inquiries and businesses trying to figure out how to do this," he said. ■

Diabetes app leads the way for open source in health

Glucosio's creator envisions crowdsourced medical apps

BY MADISON MOORE

After being diagnosed with Type 2 diabetes, Benjamin Kerensa, a 31-year-old software developer, searched for an app that would track his glucose levels and metrics.

He searched, and didn't find much. So, he created his own app: Glucosio. It is the only open-source diabetes app that does glucose tracking with third-party integrations and crowdsourced research, led by Kerensa himself. It gives open-source developers the ability to use, copy, study or change the

As a long-term goal, Kerensa wants to turn the assistant into something similar to Google Now, but for diabetics, and without storing all of the user's data in the cloud. He said that Glucosio respects the privacy of users and allows them to be in control of their data.

He said he is continuing to add features that will help people with diabetes better understand their disease. One thing he is excited about is the possibility of integrating with platforms like openEHR to create built-in medication and glucose reminders, or to

that will play out in health and medicine," said Kerensa.

How Kerensa got here

While the future for him seems to be a focus on software innovation, his past is filled with other notable open-source contributions.

He started with open source in 2001, when he installed his first Linux distribution, and at the time, he was an early user of Mozilla Firefox. He contributed to Ubuntu, an open-source Linux distribution, and also to projects like Debian, Ganeti, GNOME. He started getting involved with Mozilla in 2009, and for two years he was a Firefox Community Release Manager, overseeing the release of Firefox Nightly and Firefox ESR.

Through all of his experiences with open source and working on Glucosio, one thing he urged developers and open-source projects is to understand the importance of privacy. He said that since working in the open is transparent, it allows developers to trust the software being developed by a project or individual developer.

"I feel like we have duty to try to ensure to the best of our abilities that the software we are making is serving the end user," said Kerensa. "If privacy is compromised, it really cannot be serving the user best, can it?"

Kerensa has taken a break from contributing since his main focus is now with Glucosio. He thinks that open source has the potential to bring big changes to healthcare, which is why he wanted to develop an open-source app that helps people impacted by a disease. With more developers and enterprises getting involved, he thinks healthcare could be "the next big thing in open source."

"If we can bring Linux to cars and phones and IoT, then there is no limit to what we can do with open source," said Kerensa. ■

'Because of how common diabetes is worldwide, pretty much every contributor to the project knows of someone affected by diabetes.'

— Benjamin Kerensa



source code as a way to contribute to the project. And contributors can find the source code on GitHub.

"We are seeing new users install the app on a near daily basis, and the open-source project has been gaining a lot of attention in the greater open-source community," said Kerensa. "Because of how common diabetes is as a disease worldwide, pretty much every contributor who participates in the project knows of someone affected by diabetes."

Currently, Glucosio allows a user to track blood glucose levels, anonymously support diabetes research by contributing demographic and glucose trends and get helpful tips through an "assistant" with Glucosio. The assistant is currently not a live assistant; it gives predefined advice for diabetes like staying hydrated or checking blood sugar.

add Optical Character Recognition (OCR) so people with diabetes can scan the display of their glucose meters and input their last reading displayed on the Glucosio meter.

The initial goal of Glucosio was to allow opt-in-based crowdsourcing of diabetes data for researchers, and Kerensa said the goal for this year is to ship its API and "turn on that feature in the back end for those who have opted in and begin collecting anonymized glucose and demographic data that we can share with researchers." Some of this data includes the glucose and activity levels of people of different genders, ethnicities and age groups. The more data they collect from users, the more researchers can use to conduct further studies.

"I'm very excited about the future of Glucosio, but most importantly about the future [of] open source and how

Building Your Own



Solution?

A patchwork ALM solution introduces multiple points of failure, multiple versions of the same information, and no one to help when things fall apart.

Seapine ALM is a single solution for managing all product development phases and artifacts from requirements and user stories through QA cycles and deployment.

With **Seapine ALM**:

- > Know you have **accurate information** with all development data stored in one place.
- > Get work done fast with **blazing performance**—even on the largest projects.
- > Develop the way you want to, whether **Agile, Waterfall, or a combination**.

It's time to upgrade from your patchwork ALM solution to a fast, reliable, single source of truth.

It's time for Seapine ALM.

Learn more at www.seapine.com.



Mark Your Calendar!



InterDrone

The International Drone Conference and Exposition

September 7-9, 2016

Paris Hotel, Las Vegas

***More than 125 exhibitors!
Demos! Panels! Keynotes!***



A **BZ Media** Event

**"It looks like the drone industry
has chosen their go-to event!"**

—Robert Rodriguez, President of the
Society of Aerial Cinematography



InterDrone is Three Awesome Conferences:

Drone
TECHCON

For Builders

More than 35 classes, tutorials and panels for hardware and embedded engineers, designers and software developers building commercial drones and the software that controls them.

Drone
FLYER

For Flyers and Buyers

More than 35 tutorials and classes on drone operations, flying tips and tricks, range, navigation, payloads, stability, avoiding crashes, power, environmental considerations, which drone is for you, and more!

Drone
BUSINESS

For Enterprise and Aerial Drone Service Businesses

Classes will focus on running a drone business, the latest FAA requirements and restrictions, supporting and educating drone buyers, marketing drone services, and where the next hot opportunities are likely to be!



**The Largest Commercial
Drone Show in the World**

www.InterDrone.com

**Registration Opens
February 1, 2016**

Android Studio 2.0 previews a shift in focus

New Instant Run, Android Emulator address need for speed, testing

BY DAVID RUBINSTEIN

Google has announced the preview release of Android Studio, the company's developer environment for mobile applications that focuses on speed of delivery and testing Android apps for every device.

Stephanie Saad Cuthbertson, group product manager for Android Studio, outlined three key additions in the preview. A feature called Instant Run lets developers see the impact of changes to their code in an emulator running the application. "Developers can now see changes in one to two seconds," she explained. "As a developer, the best feeling is when they're in the flow, and by letting them see the results of their changes that quickly, it helps them stay in that zone. Developers tell me when it comes to mobile app testing, emulators are critical."

The second major feature is Android Emulator. Cuthbertson noted there are 1.4 billion 30-day active devices in use today, and organizations "want to test on all of them." A rebuilt user interface presents common device actions such



Android Studio 2.0's Instant Run feature previews changes to code in real time.

as volume control and screen rotation, and in the new version, developers can rescale the emulator window, she said. "In Android Studio, you can build a matrix of system images and emulated hardware" to test against, she added.

Android Studio 2.0 also includes what Cuthbertson called "a very early preview" of a new GPU Profiler that will enable developers to record and replay graphics-intensive applications frame by frame for debugging as the application runs the OpenGL ES Android code.

Cuthbertson discussed the "tremen-

dous momentum" Android Studio enjoys, with five major releases since it came out of beta last year. Versions 1.2 and 1.3 of Android Studio reflected Google's investment in "stability and quality," she said. "Developers spend hours at a time in their environment; it has to be amazingly stable."

Android Studio 2.0 is built on the IntelliJ Java development environment created by JetBrains, and the speed of updates to that tool is driving the work in Android Studio, Cuthbertson said. With the new Test Runner tool from IntelliJ, Android Studio users can now do unit tests and Android tests at the same time.

Further, IntelliJ's approach to "intention action"—suggesting a solution when classes haven't yet been imported or methods haven't yet been written—is behind the Android Studio ability to test deep links in Android app, said Cuthbertson. "In the IDE, you'll see a small bulb that suggests improvements to code. They're smart suggestions in context, like, 'Hey, you should add a link here,'" she explained. "You can do static analysis on links, and real-time testing." ■

NodeOS nears 1.0

100% JavaScript OS moves closer to reality

BY ALEX HANDY

The Node.js community has built itself an operating system. Known as NodeOS, this Linux kernel-based distribution layers a 100% JavaScript world on top of the traditional OS foundation. It is backed by the Node Package Manager (NPM), which allows JavaScript applications to be downloaded and installed easily.

NodeOS is a slender operating system targeted at server needs. The project lives on GitHub, but is also available as a compiled image that is rebuilt after every commit.

While NodeOS is an exciting next step for the Node community, the project is still lacking in full support for tra-

ditional computing. Though NodeOS supports Linux Containers and Docker, it does not support mounting file systems within those containers. And file systems are a major sticking point for the operating system, as only RAM disks are supported at present.

That means file systems are destroyed at shutdown. User space does have access to home folders, which can be presented to the operating system at boot and then associated with a specific user. Root will receive this treatment first, and will have access to all potential home folders. Users will also be able to upgrade home directories for a user, later designated as Root.

Each layer of the operating system has its own future planned as far as file systems go. The initramfs archive is used to create the temporary RAM disks used to run a NodeOS machine. The rootfs will be read-only, and users will behave as a traditional userspace file system.

The goal of NodeOS is to build an ephemeral operating system for quickly deploying Node.js applications without requiring separate images for every application. The disposable nature of NodeOS means that configuring an instance will inherently be part of booting that instance, making provisioning and configuration tools essential.

NodeOS will continue to focus on NPM as the primary method of software installation. With the exception of the Linux kernel, the goal is to have the OS be made up almost entirely of JavaScript. ■

Data Made Simple


RavenDB replication can be setup in less than five minutes, allowing a single database to distribute all changes to an unlimited number of nodes.

Customers are using RavenDB replication to deliver high performance (through load balancing) and enhanced availability.

The RavenDB client is self-learning and independently maps out the replication topology. If there is an error, it will automatically failover to a secondary node until the problem has been resolved. You don't need to have a babysitter watching anymore - as a second generation database, RavenDB just knows how to figure things out.

www.ravendb.net

Proudly Developed by



HUD wants to put high-speed Web into low-income housing

Julián Castro advocates for public, private effort for broadband

BY ALEX HANDY

The “digital divide” may seem like a foreign concept in a world where even children have a smartphone, but just having a phone is not enough, said U.S. Secretary of Housing and Urban Development Julián Castro. He was in Silicon Valley to discuss the ConnectHome program: an effort to bring broadband access to 750,000 low-income American families living in public housing.

“This is fundamentally about sparking greater opportunity in the lives of families of modest means and ensuring that 21st century America remains the undisputed land of opportunity for all people,” said Castro, speaking at GitHub’s headquarters.

“What we’ve found is that more than half of low-income families are not connected to the Internet, so ConnectHome [is making] sure the Internet follows young people home, and that all families are able to avail themselves of 21st century tools in this 21st century economy,” he added.

Castro was joined by GitHub’s CEO and cofounder Chris Wanstrath, and Chike Aguh, chief programs officer for EveryoneOn. Aguh’s non-profit specializes in connecting low-income families with discounted broadband on city-by-city basis.

Aguh said that EveryoneOn is working with HUD to ensure kids aren’t limited to phones or school for Internet access. “We work with national partners to make sure we have Internet services that are affordable to families. Then we work on the ground with communities. We come to this because we believe no matter who you are, you should have access to the Internet and the opportunities it brings,” he said.

While being the Secretary of Housing and Urban Development may sound

like a powerful position, Castro said his office is limited in the way it can encourage broadband service discounts. In this regard, he said that partnering with non-profits and for-profits is essential.

“We’re trying to...bring forward a partnership approach [and] move straight forward; EveryoneOn are the key to this,” said Castro. “The fact is that in implementing ConnectHome, we are limited in terms of what we can



do. It takes third-party non-profits to bring the parties together.”

For GitHub, bridging the digital divide means offering more than 500,000 free private accounts to low-income families. Wanstrath said he hopes to hire from the ConnectHome user base in the next five years. “The key is, if you want to be a programmer, you need the Internet,” he said.

“You need a computer. There are still way too many kids who don’t have access in this country. We have employees, customers and users who’ve had amazing opportunities, and we want to help give that opportunity to as many people as we can, even in our own backyard.”

Wanstrath went on to promise his employees would contribute more than 2,000 hours of volunteer time over the next two years, teaching programming

and other technical skills to the communities they live in. Castro praised GitHub for its commitments, which Wanstrath doubled on the day of the secretary’s visit.

“There is a conversation raging right now about the demographics of the workforce here. How do we ensure we have a Silicon Valley workforce that looks like the world?” said Castro. “How do we start including people from modest backgrounds and people of color? One way is to hire people who have lived in public housing, and many of them are people of color. If they have the chance to get savvy about the Internet and develop more technology skills to get them into this industry, that’s a win for everybody. That’s

why what GitHub are doing makes a lot of sense.”

But Castro also pointed out that more help was needed to continue to bridge the divide; most unconnected families are also without a computer.

“One of the needs we see clearly is for more devices,” said Castro. “That’s one of the things ConnectHome is looking for: more help with devices. We’ve identified that as a very particular need. A lot of people have a flip phone or smartphone, but that’s not something you do your homework on. We need devices that can provide a fuller, more meaningful experience.”

Aguh suggested that for US\$2 million, all of the homes targeted with ConnectHome could be outfitted with a \$90 device. He asked for Valley companies to help raise those funds.

To date, ConnectHome has brought Internet into the homes of between 175,000 and 200,000 publicly housed families. The end target for the project is to connect 750,000 homes within the next two years. ■

We know how application development can be.



Let DotImage take some of the bite out of your challenges.

Atalasoftware
DotImage®

Connecting the dots is a no-brainer. DotImage image-enables your .NET-based web application faster, more cost effectively, and less painfully than if done on your own. This proven SDK is versatile, with options including OCR capabilities, WingScan compatibility, and support for a range of formats. Coupled with dedicated assistance from our highly knowledgeable and skilled engineers, DotImage helps your business connect with powerful information hidden inside your documents, making the big picture much easier to see.

SYNCFUSION SUCCINCTLY SERIES



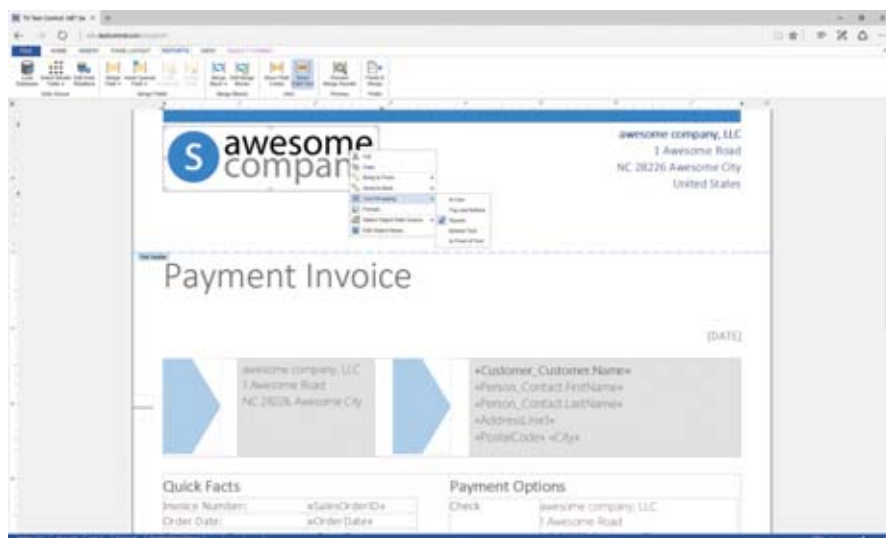
80 titles and growing | Ad-free | 100 pages | PDF and Kindle formats

DOWNLOAD YOUR FREE COPY TODAY!

syncfusion.com/sdtspark



WINDOWS DEVELOPER WATCH



MS Word-compatible drawing and shapes features are now available in TX Text Control.

TX Text Control X13 reaches out to ASP.NET MVC framework

BY CHRISTINA MULLIGAN

Text Control has announced the latest version of TX Text Control for all platforms, including ASP.NET, Windows Forms, WPF and ActiveX. The latest version features a request from users to include an MVC version of the ASP.NET product.

"We are very happy about this release; it was a lot of work," said Björn Meyer, president of Text Control. "The core focus was on ASP.NET MVC as we received so many requests for this. And finally, we can deliver a fully featured MVC version."

The ASP.NET product features MVC versions of its HTML5-based WYSIWYG rich text document editor and reporting template designer. With this version, ASP.NET developers can integrate a fully featured MS Word-compatible word-processing interface to their cross-browser Web apps without the need for any plug-ins or client-side footprint, Meyer explained. The required client-side MVC references and files are available through a NuGet package.

In addition, MS Word-compatible drawings and shapes have been introduced to all TX Text Control .NET products. Users can add shapes as single objects into TX Text Control, and completely edit them using the ribbon interface, a ready-to-use dialog box, or programmatically with the TX Text Control API.

The reporting framework also got a new merge block concept to make handling inside the template designer more intuitive and fault-tolerant. Merge blocks can be used to repeat content bound to a data table, and help non-developers create templates using only MS Word skills.

Other features include Unicode bullet characters, extended document properties, a performance boost, and new document viewer in the ASP.NET, ActiveX and WPF versions of TX Text Control. There are also undo grouping, new dialog boxes, and a built-in mouse interface for inserting text frames in the ActiveX, Windows Forms and WPF versions. ■

In other news...

■ **Atalasoft**, a provider of imaging tools, has announced an update to its .NET and MobileImage SDKs. Document Imaging SDK 10.6.1.2 aims to resolve problems from its previous release. Fixes include an improvement to its OfficeAdapterDecoder, a fix to its AutoColorCommand, and fixed thumbnail overlap in certain conditions when using Internet Explorer 8. MobileImage 2.5.0 was released to fix an out of memory error when using the barcode reader, an incorrect BarcodeReader direction map, and overlapping messages in AnimatedDocumentCaptureExperience.

■ Software development tool provider **DevExpress** announced an update to its enterprise-ready development tool set for desktops, Web and mobile. DevExpress Universal version 15.2 features new tools to help developers create high-performance on-site or cloud-based solutions targeting Windows, the Web, iOS and Android. Features include a new diagram control, report designer, form layout widget, data grid widget, HTML5 report viewer, formatted data bars, adaptive data grid, and a query builder.

■ Provider of document, medical, multimedia, raster and vector imaging toolkits **LEADTOOLS** has announced major updates to its document and medical product lines. Enhancements include improved document formats, updated DICOM specification compliance, and new features to the Medical viewer control on multiple development platforms.

■ **Syncfusion** has announced a new end-to-end solution for creating and sharing interactive business dashboards. The company announced the new Dashboard Platform, which allows users to compose dashboards using a drag-and-drop designer and rich widgets for visualizing information. It also integrates with the Syncfusion Big Data platform, and it supports Microsoft Excel, CSV, Microsoft SQL Server and Apache Spark. ■

2015:

THE YEAR

CONTAINED





BY ALEX HANDY

Containers are spreading through the software industry like wildfire. Companies high and low are supporting Docker: Microsoft, IBM, HP, Oracle and VMware, and many other companies have added container support, compatibility, management capabilities, or at the very least are planning to release something this year.

Why so much fuss over containers? They use fewer resources than virtual machines, but we're hearing more and more that the idea of having a single object that is the deployable is the end-game appeal.

Containers may just spark an uprising in the world of hosting, too. Naturally, you can host containers in Amazon, but Google and Joyent are offering container-based hosting as a first-class citizen, with management tools that can ease the pain of deploying thousands of containers at once.

This could mark a turning point for those cloud hosts, as a window opens to support workloads that Amazon Web Services might be a little less efficient at handling. Google, in particular, is doubling down on its hosting services with new integrations across developer tools, as well as a startup accelerator aimed at building new businesses in its container-based systems.

And yet, there are still concerns around the ecosystem that has risen up around containers. Primarily, the worry is that Docker is not addressing all of the needs arising from enterprises moving such systems into production. While Docker 1.9 moved the dial quite a bit in terms of making Docker enterprise production-ready, the company still has security, packaging and management issues to figure out.

For now, Docker seems to be the king of the container game, but 2016 will be the year we see the platform first starting to deal with competition. CoreOS has quickly been ramping up its rkt (Rocket) container platform. As a container-focused OS company, CoreOS has quickly taken up the banner of giving the market an alternative to Docker. 2016 will be the year we start to see the fruits of their uprising.

For the average enterprise, however, 2016 is likely to be the year containers start making their way out of test and research departments, and begin appearing on customer-facing servers in production. While some of the core elements may still need work (like Docker's registry and packaging tools), you can bet that those tools will be shaping up in order to make the case for production-grade containers before 2016 ends.

Continue reading for the rest of the 2015 Year in Review. ■



Microservices are bringing SOA back

BY ALEX HANDY

Could it be that the SOA revolution of 2005 was simply premature? After years of pain and suffering, it would appear that the Service-Oriented Architecture is, in fact, all the rage. What our modern SOA landscape looks like versus what it looked like in 2005, however, shows quite a land shift.

Ten years ago, services were hosted as a single application on a single server, and were likely using SOAP and XML. Today, services are hosted by the hundreds on virtual servers and inside containers. SOAP long ago gave way to REST, and while XML is still everywhere, JSON has risen quickly with the JavaScript crowd.

That crowd has also been a big part of why services are finally giving way to microservices. The Node crowd loves its microservices, and they've designed solutions for hosting, managing and deploying them. NodeOS, in fact, aims to be an environment for hosting just such a retinue of microservices in a simple, package-managed fashion.

All this, however, means the good ol' SOA tools must be reinvented and redeployed. For some, like Uber, this means designing their own system based on HAProxy and their open-source Ring Pop suite of node-management status tools. Indeed, HAProxy has sort of run away with the service registry market for startups, and enterprises are starting to realize its power as well.

It's not going to last forever, however. The last time we did the services thing, registries, repositories and management systems were the lifeblood of dozens of startup software companies. Many were sold, some failed, and even some still remain in business today. But the unfortunate truth is that the world in which those solutions matured is long gone.

As Uber has demonstrated, the future is filled with hundreds, if not thousands, of microservices, all hosted and running on the same registries, repositories, infrastructure, and datacenter operating systems. That's why the hot new startups are all building their own registries, and why things like GitHub LFS, Apache Archiva, JFrog's Artifactory, and Maven's Nexus are growing inside enterprises.

The microservices revolution has already taken place. Now, we've just got to figure out how to deal with it as services propagate like kudzu around our datacenters. ■

Agile was the

Organizations realized its benefits at the team level and worked to scale it to the enterprise level

BY CHRISTINA MULLIGAN

Agile has been a sought-after process for the majority of companies in the software development industry for many years, but 2015 appeared to be the year they all finally started to understand and see the benefits.

"Agile is quickly becoming the de facto standard for software development, and five to 10 years from now it is going to go beyond the de facto standard and become the best practice and the only way to do software development," said Robert Holler, CEO of VersionOne.

The Internet of Things

BY CHRISTINA MULLIGAN

Much ink was spilled in 2015 about the Internet of Things. It's a concept that has been bandied about for several years—remember the use case of your refrigerator letting your local grocer know when you're out of milk? The coming together of smartphones, sensors, GPS, Bluetooth and other technologies that enable mobile interconnectivity has moved us closer than ever to a reality of ubiquitous interconnectivity.

In the beginning of the year, Intel's Open Interconnect Consortium (OIC) unveiled a preview of IoTivity, an open-source software framework designed to provide connectivity for IoT devices. The Online Trust Alliance followed suit with the release of its Internet of Things Trust Framework in August, designed to address security.

IBM made efforts to help developers create IoT solutions in July when it released its developerWorks Recipe site, which provided templates and tutorials for building IoT apps and the intelligent logic behind the devices. Salesforce entered the IoT market in Sep-

new norm in 2015

With companies recognizing the benefits at the team level, they strived to achieve the same benefits at the enterprise level, and solutions were there to help bridge that gap. QASymphony updated qTest eXplorer in the beginning of the year to help organizations take agile software testing to the enterprise. Mendix and Pivotal teamed up in the middle of the year to power modern application development for agile enterprises. VersionOne focused on best practices for agile in its 2015 Spring Release, providing new features to give organizations better insights into their agile initiatives.

In August, the Agile2015 conference focused on the idea of scaling agile. AgileCraft unveiled the Team Reconcil-


iation and Extraction engine, designed to normalize data for agile organizations at the program, portfolio and enterprise level. Tasktop updated its Sync and Data solutions to connect everyone in the software life cycle; and ThoughtWorks unveiled an enterprise agile solution, Mingle Plus 2.0, that aimed to bring the value of agile to the enterprise without sacrificing visibility, flexibility or team productivity.

Going hand in hand with agile, organizations embraced the DevOps movement to create the environment they needed to keep up with agile. While DevOps isn't a new idea in the software development industry, the DevOps Enterprise Summit in Octo-

ber highlighted the fact that the method was just getting started in the enterprise. Solution providers emphasized that as they looked to ease the culture transformation throughout the year.

Kony and SOASTA partnered in the beginning of the year to provide new integrated mobile DevOps solutions that brought agility, usability, certainty and efficiency to teams. In May, Microsoft announced several new DevOps solutions to help enterprises identify best practices and avoid common pitfalls.

The year ended with OpsClarity looking to tackling some of the Ops problems in DevOps. The company released a new platform that brought machine learning, Big Data analytics, and data science to DevOps teams as a way for them to process and visualize massive amounts of data in real time. ■



tember with the announcement of its IoT Cloud, which allows developers to push their IoT device data into the cloud. Verizon turned its focus to the Internet of Things in October with a global strategy designed to make developing IoT solutions easier. The strategy included the launch of ThingSpace: an IoT platform for developers, a new dedicated network core and connectivity options, and three new end-to-end smart cities solutions.

And it wasn't just IoT that made advances in 2015; artificial intelligence solutions became smarter and smarter as the year went on. In January, researchers from the University of

Alberta said they made a breakthrough in AI when they created what they believed to be the first program that solved a non-trivial human game with imperfect information: heads-up limit Texas hold 'em. Another game-playing system was announced in September when a computer engineer from Imperial College London developed an AI system that used machine learning to teach itself chess.

And IBM's Watson computer used machine learning and artificial intelligence to drive innovations in health-care, banking and even in talking toys

that used natural language to communicate with children.

Toward the end of the year, Facebook announced advancements in image recognition and natural-language understanding; and Microsoft researchers developed a new system that analyzed images and interpreted them as humans would, using multi-step reasoning and deep neural networks.

Other organizations such as the FAA vowed to keep the public safe when a new industry popped up: commercial drones. Drones have been used by the military for years, but last year the market started to move into agriculture, search and rescue, surveying, photography, and videography, and with them came new tools for developers.

3D Robotics announced DroneKit, a free and open drone application development platform in March. DJI released the Matrice 100 in June to give developers a new platform to research and test their drone applications and technology, as well as Guidance to provide an object avoidance system for drones. The company also launched the Manifold in November to give developers a new embedded computer for drone development. ■

Security remains a stepchild

Despite advances in productivity, work on fixing vulnerabilities lags

BY MADISON MOORE

Software security is improving, but this past year still saw hacks and security breaches. In 2015, companies were creating new tools or initiatives to make sure data and critical information were protected, but with a fair share of leaks and hacks, the wait for a solution to software security continues.

Unlike other approaches to security, a company called SourceClear started by raising funds to improve software security in October, saying that its software works inside a developer's workflow and with a team's tools, enabling visibility into the risks of using other people's code in real time as the developers work.

Also released in October was Cigital's most recent findings of its Building Security in Maturity Model, declaring that software security is in fact lagging. With the release of this study, the application security firm added the healthcare industry to its analysis, joining financial serv-

es, independent software vendors, and electronics. Gary McGraw, CTO of Cigital, hoped that these findings would get companies to "buckle down" and focus more on security in the months to come.

October was a big month for announcements, as the government also joined in on the fight for cybersecurity. The U.S. Senate passed a controversial cybersecurity bill known as The Cybersecurity Information Sharing Act (CISA) in October, and if the bill is signed into law, it would allow businesses and government agencies to share information related to hackers and their methods. Several organizations like Twitter, Yelp and Reddit, spent months trying to raise awareness about it and why it shouldn't be passed.

Other companies were busy handling their own security issues, like Dell when it had to respond to concerns about a certificate called eDellRoot that was supposed to make things fast

and easy for customers, but instead introduced a hole in security.

Instead of handling security issues, Docker was busy improving security. It had a lot of changes, starting in August when it introduced Docker Content Trust, which uses digital signatures to secure Dockerized content. In October, CoreOS and Docker, along with a group of industry leaders, wanted to create common standards for software containers through the Open Container Project, which included making sure they had a well-designed software container specification that was secure across all platforms. And, in November, Docker announced new security enhancements that safeguard and protect Dockerized distributed applications, without impacting the developer's workflow.

To sum it up, Verizon released its 2015 Data Breach Investigations Report, which revealed that while cyber threats are getting more sophisti-

Testing catches up to DevOps

BY MADISON MOORE

With mobile, agile, Continuous Delivery and DevOps becoming more popular, testing needed to catch up in 2015. Testers evolved this past year from working solo in silos, to joining developers to deliver code faster. Testers have become key components in the world of agile over the year, and they have been trying to maintain the pace of software delivery and testing constantly throughout the life cycle.

Mobile also impacted the testing community last year. Agile has changed the rate at which new application releases and updates are generated, so testing must be done against the mobile devices that are continuously being introduced to the market. Several companies, updated and released tools to test mobile applications faster and easier.

The word heard at testing confer-

ences and from test tool providers is "automation." But, as LogiGear founder Hung Nguyen noted, the growth in automation testing will position companies to offer result-type SLAs. "It's still a challenge, though," he said. "There are no standards on what a test case is, no agreement on a unit of measurement, so results-based [SLA] is very difficult."

Perfecto Mobile announced new integrations in November for mobile testing with Microsoft Visual Studio and Visual Studio Online. Perfecto's integrations connect its cloud-based test lab, Continuous Quality Lab, to Microsoft's developer tools to enable automated mobile Web and app testing via the cloud.

In December, Keynote announced the integration of Keynote Mobile Testing with Microsoft Visual Studio Team Service and Team Foundation Server.

This allowed Microsoft customers to test their applications on the industry's largest library of interactive mobile devices from the Keynote DeviceAnywhere Cloud, and to collaborate seamlessly with development teams throughout the application development life cycle, on-premise or in the cloud.

Some companies released tools to easily scale agile testing, such as SmartBear's new version of TestComplete, which supports Python, something QA and development professionals use frequently. In August, SmartBear updated its API tools to accelerate development and testing processes in agile teams.

Experts suggested that agile development practices and Continuous Delivery helps organizations push applications several times a day, making testing an afterthought. In the new year, the focus will be on making sure those silos stay

cated, many cyber attacks still rely on decades-old techniques.

That being said, in June, cybersecurity firm Kaspersky Lab announced that it had experienced an advanced and stealthy attack on its own internal networks. Sony had to settle after its systems suffered a breach in November from hackers whom the company claimed were angry about the movie "The Interview." That breach led to the release of personal data, and former employees say it happened due to company negligence. Toy giant VTech also saw a breach in November, which caused hackers to access 6 million children's information. Experts say it was due to a lack of common steps to protect passwords.

The age-old battle of software security continues, and mobile applications pose even more problems for both the developers and the applications themselves. Experts say that the responsibility of securing mobile apps shouldn't be on just the developers. Instead, security should be a coordinated effort between the business and development teams, and this is something to consider moving into 2016. ■



The year in mobile

Developers see more choices for building apps as devices move closer to replacing desktops

BY ALEX HANDY

The smartphone reached eight years old in 2015, and looking at the app stores on iOS and Android, it's becoming clear that these devices are replacing desktops. From Microsoft Office reaching both platforms, to Apple's introduction of the professional-grade iPad, smart devices became first-class work citizens.

And that's not surprising, considering the trajectory of these platforms during their eight-year histories. Today, one can find business applications galore in the App stores. And just about any language can be compiled into an Android asset, from Fortran to C#.

In 2015, Microsoft brought Windows 10 to desktops and tablets alike. Then they announced tools for building Android applications inside Visual Studio. Combine these moves with the rapid growth of Xamarin's C# development solutions for mobile devices, and you've got a recipe for a year dominated by cross-platform offerings from Microsoft.

Android hit version 6, iOS reached version 9, and both platforms saw the release of thousands of new applications. Voice recognition became the "it" feature, as teenage girls adopted the feature in droves, making it a must-have for next-generation applications.

Google updated the IntelliJ-based Android Developer Studio throughout the year, capping 2015 off with the 2.0 release. This new version offers faster builds, and before the end of the year, should also have a much faster, updated Emulator 2.0.

Apple, on the other hand, gave developers two new methods of UI interaction to think about. When the company introduced new iPads in September, it also introduced the forceful tap method of interaction. This was coupled with the Apple Pencil, a stylus designed to give users varying levels of pressure for interaction.

As always, these interface paradigms will undoubtedly filter down into other mobile platforms as 2016 bears down upon us, just as the original iPhone's interface design was quickly adopted on other platforms.

Hosting mobile applications is a big business as well. It's easy to draw a parallel between the growth of cloud hosting and the growth of mobile applications. In 2015, however, the Internet of Things began to pop up all over the landscape, adding complexity to architectures everywhere.

All those things are still a mess of control platforms, design software, and operating systems. We're hoping that 2016 will bring more coherence to the Internet of Things, as devices spread out across the landscape like so many cattle. The unfortunate fact about the IoT and the cloud, however, is that only the server-side can be culled when things go wrong. ■



torn down and testers use a unified test-automation platform. That way, tests can run in parallel, and testing can finally catch up to agile and DevOps.

With release cycles getting shorter and more devices being produced with more platforms to run, the amount of software issues that could affect a user is much higher than it ever was before. So, while testers might have more developer-like traits, it's important that their skills as a tester don't diminish. ■

WANTED

DEVELOPERS, IT PROS, INFORMATION WORKERS
& PROJECT MANAGERS SEEKING
THE BEST SHAREPOINT TRAINING!

COME TO



SPTechCon

The SharePoint
Technology Conference

FEBRUARY 21-24, 2016

AUSTIN, TEXAS

www.sptechcon.com



"No matter what your role, whether developer, IT Pro, Power User or project manager, you will find plenty of sessions to meet your needs."

— Kathy Mincey, Project Manager/Trainer, FHI 360

"If you want a great mix of awareness-building, training, all in good company and in a great location, SPTechCon is where you'll get the most bang for your buck!"

— Mary Panke, Strategic Knowledge Director, PAI

Learn what's new in SharePoint and Office 365! Look into the future with SharePoint 2016!

Whether you want to learn about what's coming in SharePoint 2016, are still making the most out of SharePoint 2013 or even 2010, or getting started with Office 365, you will find the SharePoint and Office 365 training you need at SPTechCon.

SharePoint 2016 training!

- SharePoint 2016 Design and Implementation
- SharePoint Farm Architecture and Performance: Testing on 2013 and Planning for 2016
- What's New for Power Users in SharePoint 2016?
- What's New for the BI workloads in SharePoint 2016?

...and much, much more!

Office 365 training!

- Demystifying Office 365 Administration
- The 10 Best Office 365 Features You've Never Used (But Should)
- Making Office 365 Work for the Business: Building Powerful No-Code Solutions
- How Office 365 is Changing the Face of Collaboration and Communication

...and much, much more!

SharePoint 2013 training!

- Going Mobile with SharePoint 2013
- SharePoint Designer 2013 Workflows – An Introduction
- Upgrading to SharePoint 2013
- SharePoint 2013: Online vs. On-Premises

...and much, much more!

Still on SharePoint 2010?

There's plenty of excellent information for you as well.

Check out a larger set of class listings on www.sptechcon.com.

**REGISTER BY JAN.15
SAVE \$400!**





Advances in in-vehicle software mean cars will be able to steer themselves, but how soon will human drivers be taken out of the picture?

BY MADISON MOORE

Software for cars has been increasing in complexity over the past few years. The amount of code in vehicles can reach a hundred million lines, and for the new and fast-rising self-driving vehicles, it could mean even more hundreds of million lines of code. These cars are not programmed in an “if-then” computer algorithm; instead, they rely on machine learning and pattern recognition. For them to drive continuously and without human assistance, experts said it will take 10, maybe even 20 years, and the end of driving for humans is not quite in view.

Regulators’ top priority is to make sure that these cars are safe to be on the roads, which could explain the stall in certain states like California, where some companies are already pushing for self-driving cars to hit the road by 2016. According to the state’s DMV, the regulations for post-testing deployment of autonomous vehicles are still being developed, and the DMV wants to make sure the self-driving vehicles are as safe

as human drivers before giving the public access to these automobiles. Automobile fatalities decreased nearly 25% since 2004, according to the National Highway Traffic Administration (NHTSA). In 2013, 32,719 people died in car crashes, and in 2012, 33,782 people died. The NHTSA reports that this is a historic low for automobile fatalities.

Yet the goal is to come up with cars that drive themselves more safely than humans, which is why the Department of Transportation and NHTSA support the initiatives, according to a statement for a symposium by Mark R. Rosekind, administrator of NHTSA.

He said that the potential to overcome human driving flaws (such as sleepiness, inattention or recklessness) makes the technology worth pursuing. As much as driver safety, though, the NHTSA is also concerned with cybersecurity, which he said could threaten the technology.

Concerns about safety

Since the software in self-driving cars controls important components like the

Mercedes Benz autonomous concept car
at the IAA, September 2015.



The Society for Automotive Engineers (SAE) Levels of Automation

HUMAN DRIVER MONITORS THE DRIVING ENVIRONMENT

LEVEL 0

(no automation)

The full-time performance by the human driver of all aspects of driving, including enhanced warning or intervention systems

LEVEL 1

(driver assistance)

The driver assistance system advises the driver about steering, acceleration and deceleration using information about the driving environment, with the expectation that the human driver will perform all remaining aspects of driving

LEVEL 2

(partial automation)

The driver assistance system handles steering, acceleration and deceleration using information about the driving environment and with the expectation that the human driver perform all remaining aspects of driving

AUTOMATED DRIVING SYSTEM MONITORS THE DRIVING ENVIRONMENT

LEVEL 3

(conditional automation)

The car automatically handles all driving tasks, with the expectation that the human driver can intervene when requested

LEVEL 4

(high automation)

The car automatically handles all driving tasks, even if the human driver does not respond appropriately to a request to intervene

LEVEL 5

(full automation)

The car automatically handles all driving tasks and does not require human intervention

Source: cyberlaw.stanford.edu/loda

steering wheel, gas pedal and brake pedal, this means they are at risk for being hacked, according to Raj Rajkumar, professor and co-director of the General Motors-Carnegie Mellon Autonomous Driving Collaborative Research Lab. There have been a few incidents already where hackers went through the WiFi, 3G and 4G connections into the car's software, such as when two hackers made a Chrysler Jeep stop dead in the middle of a highway. Rajkumar said that companies like Tesla, GM and anyone else working on self-driving cars need to be "sensitive to security vulnerabilities."

GM is making sure it covers potential cyber threats by using a layered approach to in-vehicle security, which means they will create a series of defenses to cover gaps in protection, possibly including intrusion-detection systems, malware scanners and other tools. They are also designing many vehicle systems, which can be updated with security measures as potential threats evolve.

"If there is a security loophole that becomes an entry point for outsiders to

get into the car, remotely and wirelessly, then that poses a big problem," said Rajkumar.

Cem Saraydar, director of the Electrical and Controls Systems Research Lab at General Motors, said the company is encouraging safety and regulation be addressed at the federal level because GM believes that legislative efforts at the state level will hurt the technology, and the government can help speed up the process. GM is participating in collaborative research to find out what safety aspects of automation systems may be important.

"If it turns out that safety regulations are required to ensure safe deployment of automation systems, then we think it would make more sense for such a safety regulation to be set by a single, federal, entity versus individual states," said Saraydar.

According to National Conference of State Legislatures (NCSL), 16 states have introduced legislation related to autonomous vehicles in 2015, 12 states in 2014, nine states (and D.C.) in 2013,

and six states in 2012. Nevada was the first state to authorize the operation of autonomous vehicles in 2011, with California, Florida, Michigan, North Dakota and Tennessee and D.C. following.

California's Department of Motor Vehicles made progress late last year when it issued a draft of potential regulations for putting humans behind the wheel of autonomous vehicles, such as Google's self-driving car. Its proposal said that a licensed driver would need to take over the car if necessary, and manufacturers would have to report accidents, create security measures to keep the car's controls safe from hackers, and tell the car's operator how to use the vehicle safely. While this is a step in the right direction, this proposal hurt Google's chances of getting their cars on the road in the next year.

Another concern of both regulators and developers is making sure the cars can understand how to handle an issue it otherwise hadn't encountered before. Some ways to test this would be to cre-

continued on page 42 ►

The start of something new

BY MADISON MOORE

The idea of autonomous vehicles began as early as the 1930s, when aerodynamics industrial designer Norman Bel Geddes created General Motors' "Futurama" exhibit at the 1939 World's Fair. His exhibition displayed the future of highways, where cars would eventually drive themselves, according to the Museum of the City of New York.

Stanford University also had a first attempt at an autonomous road vehicle in the 1960s when researchers were trying to create a platform for controlling a Moon Rover from Earth, according to Stanford research. They named it the Stanford Cart, where a magnetic tape loop made it possible to vary the time delay of steering commands. It went through several upgrades over the years, and by 1979 it was able to get through a crowded room.

In the 1970s, a professor named Sadayuki Tsugawa and his colleagues at Japan's Tsukuba Mechanical Engineering Laboratory created what might be the first autonomous car. It was a route guidance system named "Comprehensive Automobile Traffic Control System," as well as a vision-based intelligent vehicle. In the 1980s, "He conducted research on autonomous navigation of the vision-based intelligent vehicle with a dead-reckoning function, as well as control of indoor mobile robots linked with infrared-based vehicle-to-vehicle communications," according to his bio found on the ITS Laboratory website.

Carnegie Mellon University would say that they triggered the biggest leap in autonomous driving. Raj Rajkumar, professor and co-director of the General Motors-Carnegie Mellon Autonomous Driving Collaborative Research Lab, believes that they did—mainly in the DARPA Urban Challenge of 2007.

"It was the turning point in the history of autonomous



DARPA Urban Challenge 2007: Boss, the Tartan Racing robot, is built on a Chevrolet Tahoe chassis. It incorporates lidar, radar and visual sensors to safely navigate urban environments.

vehicles, proving that vehicles that drive themselves are no longer just science fiction," said Rajkumar. "It was a watershed moment which made the singular point that vehicles that drive themselves are no longer a question of if but when."

There were several memorable challenges that Carnegie Mellon participated in, with the first being in 1995, where roboticists drove NavLab 5, a 1990 Pontiac Trans Sport, from Pittsburgh to Los Angeles on a trip called "No Hands Across America." The vehicle used a portable computer, a windshield-mounted camera and a GPS receiver, and it almost completed the drive fully autonomously.

Here in 2016, headlines claim there will be fully autonomous cars on the road in the next year or two. But experts argue that it's going to take much longer than consumers think. ■

◀ continued from page 41

ate a realistic urban environment, much like Mcity, where Ford is testing their self-driving vehicles for issues that they otherwise wouldn't be able to test safely in the real world.

The DMV suggests that an alert driver can react to a situation in 1.75 seconds. Because of this, Rajkumar said the self-driving vehicle must be capable of having the "ability to slow down and come to a safe stop," just as a human driver would. The "vehicle must be capable of dealing with tricky conditions by itself."

AdasWorks' CEO Laszlo Kishonti said that the current thinking is that the cars will give back the controls to the driver if it is in a situation it hasn't seen before. (AdasWorks is a company that develops artificial intelligence for cars.)

"If these cars become popular, then people will lose their ability to react, not just because they are not focused, but if you lose your experience with driving, then human drivers will become more and more dangerous," he said.

Even if the industry manages to work out the kinks of safety and security, it's another thing to get people to want to ride in self-driving cars. Since there is a younger generation growing up surrounded by new technology, it's possible that they will be more willing to embrace this new paradigm of using cars, according to Jim McBride, head of research and innovation at Ford Motors. However, he said that, in his experience, once he has put someone who was reluctant about self-driving vehicles behind the wheel, their opinion changed.

"It's hard for a lot of people to imagine what it's like because those products exist for them to try," he said. "They've never had that experience, so imagining what that experience is going to be, they have no paradigm in their mind."

Rajkumar said that the general public might have an unclear understanding about the software and technology that currently exists and is ready for the road because of the "misinformation, hype and confusion." He said that the cars on the market are not fully autonomous and that the technologies being introduced are to assist the driver that is still actively paying attention to the road. The cars that are being tested have driver-assistance software, and drivers are still expected to pay attention at all times.

"It will take many, many years for that technology to evolve and mature," said Rajkumar. "Meanwhile, the technology has to be reliable; it's going to take time."

Companies agree about the different degrees of vehicle automation, but there is disagreement about what those degrees are, according to Rajkumar. The Society for Automotive Engineers (SAE) has defined six degrees of automation, from Level 0 of no automation to Level 5 of full automation. Rajkumar added that the NHTSA has defined four levels of automation.

GM uses SAE definitions of driving automation, saying that they provide a way to describe differences between low-automation and high-automation systems. GM's Saraydar said that commercially and widely available fully autonomous vehicles are "likely a ways out because the technical challenges to handle all driving conditions as good as, or better than, a human driver is a difficult task and it will take time to develop the necessary technical solutions."

Despite the difficulties, Saraydar said that there are demonstrations of concepts of autonomous driving capabilities that are already happening, and that "This trend will intensify across the industry."

The exact year of when self-driving cars will be completely road ready are all over the place since automobile companies are making and changing their predications of when their cars will be on the market. Glimpsing at their predictions, anywhere from 2025 to 2040 is fair game for auto-driving cars on the market.

Current self-driving vehicles and the ones that are in the making rely on the sensors and image data to understand objects, and they include driver-assistance software, which means the vehicles are just assisting the drivers. Although it seems the industry is close to having driving control solely in the hands of the vehicle, it could be

A car driven by a single embedded processor

New technology comes with a price, and until these cars are able to be mass-produced, the cost is not affordable for the average person who might spend \$30,000 to \$40,000 on a new car. Google's self-driving car is running at \$75,000 because of the pricey lasers and radars that sit on top of the vehicle, and the 2015 Tesla Model S goes for about \$70,000.

A small team of engineers wants to show the world that it is possible to create a vehicle that doesn't cost thousands of dollars to implement software or hardware, and they too are joining the race, even if they are across the country.

Located in Budapest, Hungary, the group is called AdasWorks, a spinoff of the company Kishonti, which provides high-performance graphics, compute, automotive and computer-vision solutions. AdasWorks brings in high-performance embedded programming, computer vision, parallel computing, visualization, cloud computing and more to the development of semi-autonomous systems. Currently, they are working on an autonomous control system for an automobile, which only cost them a few hundred dollars.

AdasWorks said they created the first autonomous car driven by a single embedded processor. It uses vision-based information and a front-facing camera. It also uses a safety system developed by ThyssenKrupp Presta Hungary, a steering system company.


Unlike other self-driving features on other vehicles, it does not need pre-recorded 3D laser scan maps of the road or surroundings to detect lanes and turns.

"[Other systems] are good for detecting lanes, but they don't really understand the environment; they just detect," said AdasWorks' Kishonti. "They don't make very complex decisions; they just mainly brake the car or keep the car in the lane."

He added that they are looking to develop artificial intelligence or deep learning for the car, so that the single unified artificial intelligence detects and understands the objects around it. They want the car to be able to predict what will happen in the next few seconds. An example he used was how a human driver would react differently when they see a regular pedestrian on the sidewalk, versus a drunk pedestrian. They want the car to be able to react in those same ways, he said. ■



Rearview camera image processing for parking



Neural networks analyze what's on the road

decades before the software and learning capabilities of the car are completely road ready.

What's already on the road

While it might take years for cars to drive us to and from places without us paying attention, there is still a lot of software moving out that is getting close to it.

Tesla: Tesla Motors has been dominating the news the past few months, starting when the company rolled out Tesla Version 7.0, which allows Model S tools to deliver new active safety and convenience features, including a fully integrated autopilot system with four different feedback modules: a camera,

continued on page 45 ►



Grey Cherveney and his 2003 Corvette

Sports car fanatics talk self-driving

BY MADISON MOORE

Picture the classic auto show. There are rows of customized cars with their hoods popped and doors open, ready for a wandering guest to ogle at them. The owners are standing by their prized possessions, waiting for someone to ask them to rev the engine of their souped-up car.

It's easy to picture one of these hot rod enthusiasts cruising down the highway at high speeds, but is it easy to picture them owning a car with upgraded self-driving features that will most likely require them to stay in lanes, drive at a safe (and slow) speed, and perhaps even leave some of the driving up to the car itself?

At the New York Autofest, these fast car lovers had mixed opinions on the future of this technology.

An accident on black ice left his 2003 yellow corvette totaled with thousands in damage, and Grey Cherveney of Bay Shore, N.Y., said that he would have loved to have something in his car that could let him know of the road conditions or to help him stay in the lane. His car was fixed up at the show, with new Lambo doors and a custom paint job, including a special exhaust, rims, spoilers, and a dress up kit in the engine. His car can get up to 200 mph, and he was able to get up to 175 mph driving it himself.

Even though he loves to go fast, he wants to keep this car for life, and keeping it (and him) safe is crucial.

"Safety is mandatory," said Cherveney. "If these features [of self-driving] are cost effective, everyone will have it."

He said he still would like to have control of his car, but if he had something that was there just to assist him (much like

Tesla's autopilot features), he would have no problem driving such a car.

Scott Becktold of Glen Cove, N.Y., was hesitant, saying he wouldn't mind a car with self-driving features as long as it is going to take him fast because, "I'm a New Yorker, and we like to go fast."

He had his own GT500 Shelby, which he said can get up to 200 mph, but he has only gone up to 160 mph with it. He said he could see other countries adding separate lanes for self-driving vehicles, but in "New York, it's not going to happen."

Dee Budd of Hauppauge, N.Y., was polishing her Stingray C7 Corvette, and she said she is definitely against self-driving vehicles and vehicles with additional software. She said she's more old-fashioned, and doesn't want a car driving her because, "If someone is going to make a mistake, it's going to be me [not the computer]."

Her son Vinny Budd agreed because recently he had his Jeep recalled due to hacking issues. He said they were shutting down the brakes and hacking the radio, and that if people could do that, what is stopping them from driving a person into a wall or trying to kill them?

"I don't want a car driving me," said Vinny. "I like to have control."

Dee said sports cars are a passion of hers, and those are the cars she prefers. She said that because hacking and security is a big problem, she is wary of all of the software and computer components being added to cars.

"I don't think we are safe anymore," she said. ■

◀ continued from page 43

radar, ultrasonics and GPS.

Tesla's Autopilot allows the vehicle to steer within a lane, change lanes with the tap of a turn signal, and manage speed by using active, traffic-aware cruise control. The car can also scan for a parking space, alert when one is available, and parallel-park on command.

Youtubers and other car enthusiasts have posted videos of these Tesla updates in action, showing hands-free driving, the car changing lanes, and even the car appearing to dive into traffic when it's supposed to stay in the lane. Despite these videos, Tesla released its autopilot as a beta. Drivers are still supposed to be actively paying attention to the road, which is probably why the new feature also comes a small warning located at the bottom of the dashboard: "Always keep your hands on the wheel. Be prepared to take over at any time."

General Motors: Although they've been a little quieter than Tesla, General Motors is gearing up for its own self-driving vehicles. By the end of 2016, General Motors will have a fleet of self-driving Chevrolet Volts that will cruise the GM Technical Center out of suburban Detroit.

The fleet will operate within the roads of the campus, and the object is to gain an understanding of how the car operates, how the sensors work and how the computer controls work, according to Saraydar. The roads on the tech center are not public, but there are still vehicles on the roads. The fleet will be able to interact with other vehicles and pedestrians, and GM will be able to learn much about self-driving from this "learning laboratory" of vehicles.

Saraydar said that GM is anticipating the complexity of the software to grow in the near future. GM develops their own software for some of its embedded controllers as well as for active safety features. They will continue to do this in their own engineering groups as they move into automated driving.

Google: Google's driverless fleet of modified Lexus SUVs are cruising around Mountain View, Calif., and Austin. According to Google's Self Driving Car Project, its software helps classi-

fy an object, predict what the car will do next, and maneuver accordingly. While Google's fleet is on the road, their goal is to use the cars to get feedback on how they perform and all of the real-world scenarios a car might encounter.

They've capped their cars at 25 mph for "safety reasons," according to a Google blog post. But some might argue that this is too slow, such as the officer who pulled over a driverless Google car recently. Lucky for the developers, the car did not receive a ticket.

Mitsubishi: Last year, it was reported that Mitsubishi Electric is using machine-learning algorithms to analyze vehicle data and driver behavior to keep drivers alert. Katsunobu Muroi, a spokesperson for Mitsubishi, said that this is different from self-driving cars because the "goal of self-driving system is to travel without human intervention." They expect it to be installed on driver-sensing commercial vehicles around 2019, but Muroi said they do not know when the units will be installed into cars.

Mitsubishi isn't the only company trying to go above and beyond when it comes to autonomous vehicles. A big announcement by Toyota in November said they invested US\$1 billion to establish a new company, Toyota Research Institute Inc., where the focus will be on artificial intelligence and robotics. Its headquarters are in Silicon Valley, and they plan to develop technology that will make driving safer for humans instead of replacing it.

Ford: Ford recently declared that they will be the first to test an autonomous vehicle at Mcity, which is part of the University of Michigan's full-scale simulated urban environment that's designed for testing these types of vehicles. "We can construct some limited number of scenarios that we prefer to do in a closed environment," said McBride. "We can do them in a high density fashion [in Mcity]."

He said that some of the testing is done in the real world, but they can set up other possible scenarios that a car might run into in Mcity, including pedestrians jumping out into traffic, bicycles diving in front of cars, or having

an oncoming car run a red light. They will use a virtual environment, complete with virtual people, which will test the cars' ability to expect the unexpected.

Driving toward the future

As the systems become more intelligent, automobile companies and software companies are going to need more individuals with the skills and knowledge to take this innovation above and beyond the road.

This advancement isn't one that will take away jobs from individuals, and according to GM's Saraydar, it will open up endless possibilities for software developers, IT experts and other software-savvy individuals.

As software capabilities continue to increase, there will be a need for IT skills for both in-vehicle technology and also simulation capabilities. "As these technologies progress and increase in acceleration to the marketplace, we will need more skilled people that understand various aspects of automated driving, but are also skilled in the art of software," said Saraydar.

GM's Rajkumar said that when maintaining these cars, two services will be required. Hardware will need to be updated (the sensors, actuators and computing platforms) when they fail or need re-calibration due to accidents, weather or fender-benders. Software will also need to be updated, and he said ideally these updates will be over the air.

Ford's McBride said that it's important to realize there is not a "one size fits all" for autonomous vehicle. He means that there are a variety of possibilities for autonomous vehicles, like an urban taxi fleet, Amazon or FedEx delivery vehicles, personally owned and operated vehicles, or commercial vehicles.

"You can imagine there is a broad spectrum of use cases, and in each of the use cases there might be different customers or partners that would want to use [the vehicles]," said McBride. "For each of those, there would be different software applications to go along." ■

Read this story on
sdtimes.com



Dynatrace hooked me up with **application performance** metrics my dev AND ops teams buy into. Now we've turned our war room into something more..... *fun*.

We use Dynatrace to record and playback scripted transactions to eliminate inter-departmental rivalries and finger pointing. The proof is in the data. Sharing it this way improves everyone's game.

Your serve...



Learn more at:

dynatrace.com



DIGITAL PERFORMANCE MANAGEMENT

Buyers Guide

Effective APM requires more data

BY LISA MORGAN

Application performance isn't just about speed anymore; it's about enabling high-quality user experiences that deliver tangible business value. Achieving that goal requires fast, easy access to more types of data than traditional APM solutions have provided.

"You have to understand each and every transaction from the start of an action on the user interface all the way down to the back-end application or the components that make up an application," said Nicolas Robbe, CMO of Dynatrace. "It used to be that the response time of an application was a proxy for the customer experience; my application runs fast, the response time is good, and therefore, the customer experience must be good."

Inferring good customer experience is unwise when user expectations are at an all-time high and issue resolution times are expected to be at an all-time low, despite the ever-increasing application and environmental complexity.

"More companies are competing via their software because that's how customers engage with your brand," said Al Sargent, senior director of product

Solutions must provide end-to-end insight so issues can be spotted and resolved quickly

marketing at New Relic. "Your software needs to be as good as it can be across three dimensions: application performance, user experience, and the ability to drive business outcomes."

The growing scope of APM, and the growing universe of data necessary to meet the needs of Dev, DevOps, Ops and lines of business, is causing some solution providers to re-categorize their offerings. For example, New Relic is now promoting "software analytics," while Dynatrace prefers "digital performance management." Others, including Hewlett Packard Enterprise and CA, are simply expanding the capabilities of their APM offerings to meet the needs of modern organizations.

The data you need today

Delivering great user experiences can be very challenging, especially when so many potentially disruptive factors are outside a software team's control. To

minimize the effect of a disruption occurring on the front end, the back end, or anywhere in between, APM must provide insight end to end so the root cause of an issue can be identified and resolved quickly.

"To truly measure user experience, APM requires that you measure the interaction with the software, the user response time, the usability of the software, the design of the software, and also what they're saying on social media," said Joe Fox, director of product marketing, application performance management, at Hewlett Packard Enterprise. "For a mobile application, you need to be able to trace those user interactions or transactions end to end across a composite application through a mobile gateway or a Web server through to the application where the data processing is done, back into a database where there's potentially some middleware component. Then, you need to be able to distribute the transaction to additional services or a mainframe or other data store."

Componentized applications and environmental complexity dramatically increase the amount of monitoring that is necessary and the amount of data collected. To be of value, APM solutions

continued on page 49 ►

Want to win in the ever-changing app economy?

Start at the end user.

In today's world, the end user is king. CA APM solutions can help you keep end users happy right now—and adapt to whatever is coming tomorrow. You'll gain valuable insights with each customer experience, boosting your competitive advantage.

Take the first step. Visit ca.com/startattheend today!



ca
technologies

◀ continued from page 47

must be able to scale accordingly.

“Identifying the source of an application problem such as performance degradation or an outage is arguably the main use case for APM,” said Kieran Taylor, vice president of products and solutions marketing at CA. “Many don’t realize that nearly 80% of the Mean Time to Repair is tied up in locating the culprit. [If you’re] only focusing on the usual suspects, [you] may miss the true cause of an application issue.”

To avoid such blind spots, developers need code-level analysis, including session state and activity, crash analytics, stack traces and affected users. User experience teams must understand usage statistics such as startup time, spin time, session activity, and which portions of the user interface are used or not used and why. Operations teams need insight into the user experience, network and carrier performance, battery status, and network changes, as well as information about specific URLs, and information about disk, memory, and CPU usage. All of that data must be segmented by app version, platform, operating system, network, device type, and even user location, Taylor said.

“In a world of componentized apps, this means instrumenting hybrid environments where some elements such as the mainframe are in the technology team’s datacenter, but container-based app components are not,” he said. “In the latter case this means harvesting metrics specific to Dockerized applications, Node.js reporting or monitoring the operation of API gateways that are the connective tissue for many applications. Support for these modern monitoring techniques is crucial to ensure there are no blind spots when [troubleshooting] a poorly performing application.”

Some APM solutions are monitoring customer sentiment information from social networks and interconnecting with enterprise applications (such as CRM) to enable more proactive customer engagement. For example, a cosmetic company uses Dynatrace to minimize revenue leaks caused by online ordering woes. When it appears that

one of its 2 million agents is having trouble placing an order, a company representative can proactively call that particular agent to offer assistance.

“It’s a tremendous shift in how you think about performance,” said Dynatrace’s Robbe. “APM is not just a troubleshooting tool; it’s a tool the business and developers can leverage to proactively resolve issues, engage with customers, and understand the experience they’re delivering to customers.”

Trainline, the largest seller of train tickets in the U.K., is using New Relic to monitor connections to different payment gateways.

“Because they’ve stored the data in a lossless manner, they’re able to jump in



and troubleshoot the problem, understand which payment gateway is down, which transactions didn’t go through, which type of currency is being used, and which error messages and response codes [appeared],” said New Relic’s Sargent. “All that information is available in an unaggregated manner, so the engineers have the information they need to resolve the problem quickly.”

Managing the data

The amount of application-related data that can be monitored and analyzed is staggering. Modern APM solutions provide an easy way to control what data is collected, analyzed and reported, how long the data is kept, and for what purpose. For example, HPE allows its customers to set up specific summarization rules that define how long instance-level data is kept, and the granularity at which the data is summarized.

“We’re providing an end-to-end view of APM so it can be analyzed and then presented to different constituents, whether that’s business personas, developers, or technical views

for application support or operations teams,” said HPE’s Fox.

Dynatrace analytics and data visualizations simplify data navigation and accelerate root cause analysis. “We also provide the ability for people to share data and collaborate around specific issues,” said Robbe. “The collaboration aspect is critical because it provides data in the context of the users in the tooling that is most familiar to them.”

New Relic helps organizations monitor massive amounts of data at scale with confidence. “You really have to ask yourself if you want to be in the business of managing trillions of events imported into your system every day,” said Sargent. “We make sure data is available fast so our customers don’t have to worry about it.”

CA provides role- and task-relevant views of data. Its latest release allows multiple teams to collaborate so applications can be continually optimized more efficiently. Users can also filter complex datasets into time-based views that correlate performance issues with changes in an application environment.

“APM Timeline View correlates separate events such as changes to [the] network, device configurations, hardware, software or even personnel,” said CA’s Taylor. “Identifying changes and connecting them to performance issues is the largest portion of time spent on problem resolution.”

Gleaning more insights

APM solutions are interconnecting with more types of tools, services, platforms and applications than ever before to provide new insights, more relevant insights, and more granular forms of insight than were available previously.

“You need to be able to follow a transaction across multiple vendors so you can get data from all your tiers and the various underlying pieces of software that make up the experience,” said Robbe. “[That way], you have one thread, across all of your systems, that comes from different vendors and different technology stacks. [The] unified data representation enables developers to troubleshoot code in minutes over extremely complex systems.” ■

A guide to APM suites

■ **AppDynamics:** The **AppDynamics Application Intelligence Platform** provides a real-time, end-to-end view of application performance and its impact on digital customer experience, from end-user devices through the back-end ecosystem—lines of code, infrastructure, user sessions and business transactions. The platform was built to handle the most complex, heterogeneous, distributed application environments; to support rapid identification and resolution of application issues before they impact users; and to deliver real-time insights into the correlation between application and business performance.

■ **AppNeta:** AppNeta provides **APM for DevOps** and **APM for IT**. APM for DevOps allows users to monitor and optimize their applications. Features include the ability to identify trends and outliers at a glance, resolve issues faster, track user performance, and measure performance, functionality and availability from the user's perspective. APM for IT allows users to find and resolve issues. Features include end-to-end visibility and SLA tracking on any network, end-user experience monitoring with synthetics, intelligent application identification with real user data, and flexible and scalable deployment.

■ **BMC Software: TrueSight App Visibility Manager** goes beyond application performance monitoring to provide deep insight into user experience. In addition to tracking and measuring user activity at the individual or location level (on premises or off), it filters data to ensure the delivery of relevant application information without the unnecessary noise.

■ **Catchpoint Systems:** Catchpoint offers innovative, real-time analytics across its **Synthetic Monitoring** and **Real User Measurement (RUM)** tools. Both solutions work in tandem to give a clear assessment of performance, with Synthetic allowing testing from outside of data centers with expansive global nodes, and RUM allowing a clearer view of end-user experiences.

■ **Crittercism:** Crittercism's **mobile application performance-management solution** enables enterprises to accelerate their mobile business. Crittercism's solution monitors every aspect of mobile app performance, and provides a real-time global view of crash, service and mobile transac-

■ FEATURED PROVIDERS ■

■ **CA Technologies:** **CA APM** is built to be easy, proactive, intelligent and collaborative, or EPIC. It offers easy deployment of APM agents, including Node.js and PHP. Proactive identification and resolution of issues occur across physical, virtual, cloud, containerized and mobile applications. Intelligent insight comes through 360-degree mobile-to-mainframe visibility, which captures billions of critical metrics per day to verify transactions. CA APM is collaborative, aligned with DevOps methodologies to instill continuous performance improvements at every stage of the software life cycle.

■ **Dynatrace:** Dynatrace delivers developers and testers the answers they need to ensure superior user experiences, proactively solve problems, and accelerate DevOps/Continuous Delivery strategies. **Dynatrace User Experience Management** and **Dynatrace Synthetic Monitoring** help developers to proactively understand and optimize user experience; **Dynatrace Application Monitoring** is designed to maximize application performance; and **Dynatrace Data Center RUM** delivers app-aware network insights.

■ **HPE:** HPE's **Application Performance Management suite** empowers businesses to deliver exceptional user experiences by examining the impact of anomalies on applications before they affect customers. Business applications are available with proactive end-user monitoring and actionable diagnostics to measure the performance and stability of apps from a user perspective.

■ **New Relic:** New Relic's comprehensive SaaS-based **New Relic Software Analytics Cloud** provides a single powerful platform to get answers about application performance, customer experience, and business success for Web, mobile and back-end applications. New Relic delivers code-level visibility for applications in production that cross six languages—Java, .NET, Ruby, Python, PHP and Node.js—and supporting more than 70 frameworks. **New Relic Insights** is embedded in the platform, enabling customers to do detailed, ad hoc queries for real-time analytics across New Relic's **APM, Mobile, Browser** and **Synthetics** products.

tion metrics across iOS, Android, Windows Phone 8, HTML5 and hybrid apps.

■ **Dell: Foglight APM,** a smart APM SaaS-based solution, requires zero configuration. Or, monitor your applications from a single pane of glass with an on-premises solution that can be optionally combined with other Foglight performance-monitoring capabilities for database, virtualization or storage.

■ **Neotys: NeoSense** is a proactive application performance monitoring solution for all Web and mobile applications. It actively monitors the performance and availability of critical business transactions within recorded user paths. NeoSense leverages the test scenario design capabilities of NeoLoad to quickly create realistic monitoring profiles. Easy-to-read dashboards, alert triggers, and notification rules provide actionable insights for pinpointing performance issues and getting to root causes quickly.

■ **Oracle:** Oracle provides a complete end-to-end application performance manage-

ment solution for custom and Oracle applications. **Oracle Enterprise Manager** is designed for both cloud and on-premises deployments; it isolates and diagnoses problems fast, and reduces downtime, providing end-to-end visibility through real user monitoring; log monitoring; synthetic transaction monitoring; business transaction management and business metrics.

■ **SmartBear: AlertSite's** global network of more than 340 monitoring nodes helps monitor availability and performance of applications and APIs, and find issues before they hit end consumers. The Web transaction recorder **DejaClick** helps record complex user transactions and turn them into monitors, without requiring any coding.

■ **SOASTA:** The **SOASTA platform** enables digital business owners to gain continuous performance insights into their real-user experience on mobile and Web devices—in real time and at scale. ■


One source of truth

See all your data. Boost performance. Drive accountability for everyone.




Mobile Developers

End-to-end visibility,
24/7 alerting, and
crash analysis.



Front-end Developers

Deep insights into
your browser-side
app's engine.



IT Operations

Faster delivery.
Fewer bottlenecks.
More stability.



App Owners

Track engagement.
Pinpoint issues.
Optimize usability.

Move from finger-pointing blame to data-driven accountability.
Find the truth with a single source of data from multiple views.

newrelic.com/truth

Challenges in APM

BY LISA MORGAN

We asked our APM Buyers Guide sponsors about the APM challenges today's software teams are facing and how their companies are uniquely addressing those challenges. Here is what they had to say.



Al Sargent,
senior director of product
marketing, New Relic

It has never been more important to build great software, but software isn't just built in one release. Software requires constant iteration based on feedback from users.

When the APM market was created over a decade ago, it helped shed light on app performance. However, it didn't reflect how modern software is built collaboratively today by cross-functional teams, or the need to understand customer experience and the impact on the business.

Today, companies need to measure performance throughout the stack, from server to business outcomes. We call this Software Analytics. This shift from APM to Software Analytics allows companies to understand how their customers are using their apps, how the apps are impacting the business, and how the apps are behaving, by asking any question about their real-time software data.

Software Analytics implies unaggregated data to understand outliers. This leads to a massive increase in the data that companies need to manage.

New Relic enables companies to understand how a customer is interacting and experiencing their software from any dimension, through a SaaS Software Analytics cloud. The result is faster innovation, increased competitiveness and lower costs.



Kieran Taylor,
vice president of products
and solutions marketing, CA

In a world of Digital Transformation and DevOps, successful software teams must adopt lean and agile methodologies that

emphasize fast response to customer needs and business value. The challenge is that this means running smaller, more continuous releases at a dramatically faster pace. Without iterative collaboration and continuous feedback, release frequency and quality falter, negating any advantage offered by agile disciplines. APM creates continuous feedback loops from production into pre-production.

The problem is that extracting value from many APM tools requires arcane knowledge typically housed in a few experts that then become overly taxed.

The newest release of CA APM 10 is designed to protect the valuable experts in a technology organization by unlocking application insight so it is more easily shared across development and operations teams. Specifically, CA APM 10 features new, role-relevant views into what were formerly complex application maps. By removing extraneous information and focusing on what matters for the task at hand, voluminous data is converted into actionable insights that are easily leveraged by all members of development and operations teams.



Joe Fox,
director of product marketing,
application performance
monitoring, Hewlett Packard
Enterprise Software

End-user expectations are on the rise. Mobility breeds expectations of immediacy, meaning that users expect what they want when they want it. Users also expect a continuous application experience that picks up where it left off regardless of the device type or whether the user has moved from a native app to a Web app. Finally, users expect high application quality. In other words, the success of an application depends on the user experience it delivers.

Failing to monitor the user experience and failing to introduce that data earlier in the application life cycle places the business at risk.

To enable high velocity and frequent

quality releases, APM tools must be used consistently across the application life cycle. Using the same APM tool, Dev, Test, Operations, and business managers benefit from a consistent base of information and views specifically designed for each persona.

AppPulse, Hewlett Packard Enterprise's next-generation APM, uniquely addresses the needs of modern organizations. AppPulse monitors applications during test and alpha and beta releases to determine where crashes occur so development teams can drill down to the offending lines of code. Lines of Business can also monitor the application performance with an index that contemplates the entire user experience. Finally, AppPulse helps Operations troubleshoot by drilling down to the back end to determine whether a database, application server or third party is slowing down the completion of a transaction.



Nicolas Robbe,
CMO, Dynatrace

Dev and Ops groups have been transforming culturally to adapt to three challenges: accelerating the pace of release without compromising quality, taming complexity further compounded by the rise of microservices, and matching the infinitely increasing customer expectations when it comes to their digital experience.

Dynatrace offers what no other APM provider can: the opportunity for developers to collaborate more efficiently with other teams across the enterprise using a common language, shared dashboards, data, insights and perspectives.

The most mature organizations are already embracing the practice of digital performance management to fully leverage Dynatrace. As a result, these groups innovate faster, cutting through complexity and delivering better customer experiences. Dynatrace has been the longest-standing recognized leader in APM, and is now driving the evolution of APM to Digital Performance Management. ■

Read this story on
sdtimes.com



DevOps, Meet UX

It's time to make the user experience a focal point for DevOps.

Users expect the world of your apps.

Instant gratification — it's the new normal. People expect to click an application on their smart phone, their browser, or their laptop — and get what they want when they want it.

61%

expect a mobile app
to respond in less
than 4 seconds

53%

uninstall or remove
the app when it stops
responding or crashes

69%

say app issues gave them
a lower opinion of the
company that created the app

For DevOps teams, the message is clear:

it's time to make the user experience (UX) a top priority across the entire application lifecycle — from design and build to deployment to monitoring.



HPE AppPulse – Measure the Pulse of the User Experience

Continually measure the user experience from the user's perspective. Measure and monitor everything the user does and everywhere the user goes within the app — and everything that happens along the way. Identify issues having the most impact on the user experience and improve what matters most.



User Flow feature of HPE AppPulse. Know where users entered, where they exited and everything that happened to them along the way.



Learn more at hp.com/go/AppPulseMobile.com



Hewlett Packard
Enterprise



These little-known tips, hidden features, new offerings and third-party add-ons will help you get the most out of cloud development with Microsoft Azure PaaS and IaaS

BY ALEXANDRA
WEBER MORALES

Despite playing catch-up to Amazon Web Services, Microsoft Azure has quickly become a contender with its powerful Platform-as-a-Service and Infrastructure-as-a-Service offerings. With constant innovations around usability, open source and cross-platform compatibility, infrastructure management and evolving software development paradigms for new devices and applications, it can be hard to get your bearings within the vast platform.

First things first: What comprises the Azure platform? It's not turtles all the way down; in Azure principal program manager Scott Hanselman's words, the underlying layer, the "infinite hard disk in the sky," is Azure storage, where you can drill down to every virtual hard disk (VHD) image in your deployment. The next level up comprises virtual machines, which you choose, configure and manage.

On top of those VMs is a middle ground between IaaS and PaaS: Worker Roles, which are stateless cloud apps that can scale their VMs up or down. Above Worker Roles we are clearly in PaaS territory, with Web Apps, Azure Batch and HDInsight (Hadoop) for Big Data analysis. And at the top there are Web Jobs, Mobile Apps, and Media Services. These pieces are among those also available as the Azure Stack for on-premise datacenters or hybrid cloud applications.

"The first distinction to make as an Azure customer is, do I want to consume VMs, or do I want to consume the platform that Azure provides me?" said Esteban Garcia, Visual Studio ALM MVP and chief technologist for Nebbia Technology, and Azure company in Orlando. "Within PaaS, we typically do a Web app and SQL services. Those are pretty straightforward and easy to discover."

It's likely that most cloud problems you're facing have already been solved somewhere in the Azure community, advises Corey Sanders, director of pro-

gram management for Azure. "Make sure you look at the full breadth of services we offer, because they solve a lot of different problems. That's the value of having a broad platform such as Azure," he said.

Read on for tips from these Azure experts.

1. Saving pennies, saving dollars

The promise of the cloud is elasticity: sizing your deployment according to demand. Too often, Azure users are surprised by the multiple dimensions of pricing, from storage to transactions, support, bandwidth and more. The first step toward transparency in billing is not to "set it and forget it" but to monitor, measure and adjust frequently.

"One aspect of VMs that is not well known but that I think is super cool is the wide variety of VM sizes that we now offer," said Sanders. "And the fact that when we stop a VM through the portal, we actually stop billing for it. The combo of those two has resulted in very few points of contention around billing."

The Azure Billing Alert Service can create customized billing alerts for your Azure accounts, and the pricing calculator is your friend.

2. Elastic Scale for Azure SQL Database

Performance and price limits on Azure SQL Database, which offer a subset of SQL Server features, can slow you down while also burning cash.

"People use SQL Service instead of the full server," said Garcia. "Whenever you do that, you pay for Database Throughput Units. You could be paying for 10 databases with five DTUs [database throughput units], which is a small number of DTUs to use. What you can do is start using this elastic scale, say 'I'm going to assign 100 DTUs to 10 databases,' and they are going to share that processing power among all the databases." In other words, a P3 tier database costs US\$4.65 per DTU per month, while the same number of

DTUs scaled out rather than up on an S2 instance cost \$1.50 a month.

"They are all able to use that shared pool of resources rather than being constrained," said Garcia. "It gets along with the idea that cloud allows you to draw resources from anywhere as needed."

Currently in preview, Elastic Scale simplifies the scaling of data tiers from just a few to thousands of database shards via .NET client libraries and Azure service templates. High-volume OLTP, multi-tenant SaaS, and continuous data collection from telemetry and Internet of Things applications are likely use cases.

3. Preview pricing

The aforementioned Elastic Scale is just one example of many new features available at preview pricing, which may be free or 50% less than the general availability pricing. Taking advantage of preview pricing lets you play with new features, stay ahead of the technical curve, save money, and possibly beat the competition by having production-ready deployments when the features go live for all customers.

4. The Azure Portal

Also in preview is the Azure Portal, a new dashboard for accessing IaaS and PaaS deployments.

"I was not happy with the portal at first, but now it's growing on me," said Hanselman in his June 2015 TechDays UK keynote. "If you double-click on the background, you can change the theme to dark. This made me so happy."

Right-clicking on a given window pins it to the start board. Charts can be edited to show, say, CPU percentage, pricing, disk usage and more. "Don't discount the portal quite yet; it's fantastic," said Hanselman.

5. Keyboard shortcuts

Every computer user knows the mouse can be deadly, in terms of ergonomics and efficiency. The best

continued on page 56 ►

◀ continued from page 55

economy of movement is achieved with keyboard shortcuts. Launched with version 5.0 of the Azure Portal, the shortcut menu can be accessed by hitting shift + ?. Luckily, there aren't too many to memorize. You'll want to use these and more:

HUBS (LEFT MENU) SHORTCUTS:

H - show startboard

N - open Notifications hub

A - open Active Journeys hub (a Journey is the current opened group of blades; a blade is card/tab/subpage that contains some group of tiles, e.g. website properties or analytics)

/ - open Browse/Search hub

B - open Billing hub

C - open Create/New hub

CHANGING FOCUS BETWEEN BLADES SHORTCUTS:

J - move focus to the previous blade

K - move focus to the next blade

F - move focus to the first blade

L - move focus to the last blade



6. Azure resource manager

The complexity of managing websites, virtual machines and databases just got a little simpler with the addition of the resource manager in the new Azure Portal. Group and view resources (such as an instance of Application Insights along with a Web application and SQL database) as a single resource group. Deployment templates in Visual Studio are also aided by IntelliSense, which surfaces new resource providers and template language functions as you write deployment templates, avoiding pesky naming errors.

"Azure resource manager has an exciting, growing community," said Sanders. "We're seeing in templates for the resource manager—starting about six months ago—a pretty exciting pick-up in the community. We seeded GitHub with a set of these templates and put them all out fully open source. Now we have over 140 contributors and more than 200 templates available."

7. Scale Sets

Do you have Big Data or container-based workloads? You may want to

orchestrate these complex, large-scale deployments with Scale Sets. Also in public preview, Azure Virtual Machine Scale Sets let you manage and configure virtual machines as a set of identical Windows or Linux images.

"A customer can come in and say, 'I want VM Scale Sets in groups of 10, and I want to configure them all with a tool like Chef or Puppet,'" said Sanders. "The other aspect with Scale Sets that's exciting is the deep integration with Azure Insights auto-scale, which restricts cost and spending by only using the compute resources you need, responding to traffic changes."

8. Security

A Denial of Service attack can hurt your wallet, hamstringing your business and harm your customers. Like other cloud providers, Microsoft shares its security knowledge as well as build in basic protections. Azure Security Center's view lets you set policies across all your subscriptions and monitor security configurations. The good news is, the days of accidentally raising DDOS flags by testing or polling your own app are over, thanks to cutting-edge threat intelligence around malformed requests and traffic sources.

9. Site extensions

Another powerful way to add custom administration features to your Web apps is with site extensions. Write them yourself or choose from the Site Extensions Gallery. These live on the SCM (site control manager) site for administration and debugging that runs over SSL and is created with every Azure website.

10. PowerShell cmdlets

Whether you want to clean up a deployment where you have some extra VHDs and VMs lying around, provision VMs, set up cross-premises networks, or other production tasks, you'll enjoy the Unix-like scripting power of PowerShell and the new PowerShell cmdlets. As of

the November 2015 update of the Azure SDK 2.8 for Visual Studio 2013 and Visual Studio 2015, the PowerShell script for deploying Azure Resource Manager templates now works with PowerShell cmdlets. Find scripting solutions already crafted for you in the PowerShell Gallery.

11. Service fabric

Riding the microservices revolution, Service Fabric is Azure's platform for assembling cloud applications from a large collection of services. "Service fabric offers a platform that runs on Azure but also on-premises," said Sanders. "This is a platform for deploying, managing and maintaining microservices. Discovery is handled for you, and it supports stateful and stateless microservices."

12. Docker

The explosion of ways you can tinker with cloud resources, from remote desktops and SSH to portal shortcuts, has only just begun. According to Hanselman, the microservices revolution means there will soon be even more options to choose from.

"Simply stated, if I've got a tiny little 10MB PHP app sitting inside of a 5GB VHD, that's a lot of VHD, a lot of virtual machine for a small Web application," he said. "Does it really need that weight? That much security and isolation? It just needs to be in a container, and it needs to be deployable in a reliable way. Docker will provide that."

"My biggest tip and trick with Docker containers is just to deploy one," said Sanders. "If you've never done anything with Docker, there's a way to quickly deploy with a fully packaged VM and Docker in the Azure marketplace. No bringing down of the Docker engine, no pulling down the hub."

13. Azure DevTest Labs

How do you avoid using up all your MSDN credit while testing on Azure? The preview of Azure DevTest Labs lets you spin up Windows and Linux environments to deploy and test applications while avoiding cost overruns.

“With Azure Dev/Test labs, the idea is that a lot of times developers have to wait for someone to spin up labs for them,” said Nebbia’s Garcia. “This allows you to spin up environments much quicker. You can choose for it to run a maximum eight hours, and after that it gets shut down. It’s a quick way to provision environments but avoid the problem of leaving it up and running. You can push a button and have whole sandbox.”

14. Application Insights

“Application Insights allows you to dig down and find the root cause of any application issues and understand how people are using the application,” said Garcia. For example, he uses it for availability testing from different geographic locations, either as a static test that checks a single page, or as a test of dynamic application flow.

15. Kudu, CloudBerry and Senty

The Kudu open-source project is a useful troubleshooting tool and client-side process explorer for capturing memory dumps or looking at deployment. It’s also a site extension and welcomes community participation.

Another useful freeware tool is CloudBerry Explorer for Azure Blob Storage, which offers a file manager-style user interface to Azure Blob Storage.

If you’re already mucking around in the cloud, you may have e-mail update needs that can be met by Senty or similar tools. Senty was designed to work with Amazon Simple Email Service, but can be adapted for Azure as well. The savings versus a hosted e-mail solution such as MailChimp can be enormous.

16. Remote debugging

In its September 2015 white paper, “Practical Guide to Platform-as-a-Service Version 1.0,” the Cloud Standards Customer Council notes that no PaaS worth its salt should be without remote debug capabilities. “Application developers should have access to tools that enable them to control activities in

the PaaS—for example, uploading (‘pushing’) application code, binding services to applications, controlling application configuration, starting and stopping application instances,” it said.

“Such capabilities should be provided in a way that fits well with the other tools used by the developer—command-line tools, graphical tools, embedded components for development environments. Ideally these tools should work via an API that is exposed by the PaaS system—cloud service customers should look for these APIs and assure themselves that the API can be used by a variety of custom tooling code.”



Remote debugging with Visual Studio fits the bill:

Developers interact with cloud applications as if they were on-premise. Best used with Visual Studio 2013, remote debugging lets you manipulate memory, set breakpoints, and step through code—with the caveat that breaking a running process could break your live website. Save this one for pre-production sites.

17. Performance testing

Performance testing allows you to generate thousands of virtual users from around the world and test your application against the load.

“If you spin up a Web application, you’re able to do a performance test right from Azure, right in the cloud,” said Garcia. “Before, it was more on the Visual Studio side. So I can see what it looks like if 1,000 people hit my app at once. It’s very useful in knowing how to scale the application: We can have fewer servers, but make them stronger by adding this performance testing feature right within the Azure portal when you first launch an application.”

18. Easy ROI: Lift and shift

Want an instant return on your cloud investment? Eliminate idle servers that only handle periodic loads. “This is something people forget about when they’re thinking about the cloud,” said Hanselman. “Azure storage, that’s an infinite disk that’s out there. You probably have a machine sitting under

your desk and it’s got a VM, running maybe an expense reporting system. It’s something that you need to lift and shift into the cloud. There are migration tools that can help you. Literally, it’s Hyper-V in the cloud, but that’s only the most basic way of using Azure.”

19. Developer services marketplace

Before you reinvent speech recognition, check the Developer Services Marketplace for free and paid ways to extend functionality, turbocharge development, and manage cloud deployments with certified Azure tools such as Iron.io for event-driven computing, or face APIs from Project Oxford.

20. Ride the IoT wave

These days, Internet of Things projects are everywhere. Hanselman, a type 1 diabetic, movingly demonstrated in a keynote how he tracks his blood sugar and insulin pump in the cloud with Azure technologies.

Microsoft Azure IoT Hub offers SDKs, management and security solutions to harness a plethora of IoT devices, and once the data is collected, there are new machine-learning tools available to process data stored in HD Insight, Microsoft’s version of the Hadoop Big Data store. Have fun!

Start your engines

Use these 20 tips as a checklist for leveraging the vast Azure platform. The more you understand Azure, the more you see where it’s headed: “We’re seeing a blurring of IaaS vs. PaaS and starting to just see a compute platform,” said Sanders.

Redmond has clearly learned to embrace today’s polyglot cloud, and you can use that flexibility to your advantage. “Azure is not only a Windows server; they have Ubuntu, they have Linux servers, they have a new agreement with Red Hat, you can spin up an Oracle database... There’s so many different non-Microsoft technologies to choose from,” said Garcia. ■

Read this story on
sdtimes.com



We didn't invent the Internet...

...but our components help you power the apps that bring it to business.



TOOLS • COMPONENTS • ENTERPRISE ADAPTERS

- **E-Business**
AS2, EDI/X12, NAESB, OFTP ...
- **Credit Card Processing**
Authorize.Net, TSYS, FDMS ...
- **Shipping & Tracking**
FedEx, UPS, USPS ...
- **Accounting & Banking**
QuickBooks, OFX ...
- **Internet Business**
Amazon, eBay, PayPal ...
- **Internet Protocols**
FTP, SMTP, IMAP, POP, WebDav ...
- **Secure Connectivity**
SSH, SFTP, SSL, Certificates ...
- **Secure Email**
S/MIME, OpenPGP ...
- **Network Management**
SNMP, MIB, LDAP, Monitoring ...
- **Compression & Encryption**
Zip, Gzip, Jar, AES ...



The Market Leader in Internet Communications, Security, & E-Business Components

Each day, as you click around the Web or use any connected application, chances are that directly or indirectly some bits are flowing through applications that use our components, on a server, on a device, or right on your desktop. It's your code and our code working together to move data, information, and business. We give you the most robust suite of components for adding Internet Communications, Security, and E-Business Connectivity to

any application, on any platform, anywhere, and you do the rest. Since 1994, we have had one goal: to provide the very best connectivity solutions for our professional developer customers. With more than 100,000 developers worldwide using our software and millions of installations in almost every Fortune 500 and Global 2000 company, our business is to connect business, one application at a time.

connectivity
powered by 

To learn more please visit our website →

www.nsoftware.com

Code Watch

BY LARRY O'BRIEN



Larry O'Brien is a software developer who lives on the Big Island of Hawaii. Read his blog at www.knowing.net.

Let deep neural networks help you

How can programmers benefit from the “the year of Neural Nets”? Statistical machine learning techniques have been surging in popularity in academic settings for years, but 2015 was a watershed in terms of industry awareness and deployment. It was not long ago when the term “Deep Neural Networks” seemed about as dubious an explanation as “Applied Phlebotinum,” but now, Google’s open-source release of TensorFlow was the source of celebration and rewriting of business plans.

Neural networks do not do any symbolic processing and are neither “logical” nor “creative.” They are “universal approximators,” meaning that they can generate arbitrarily accurate results for a given function. The inputs are swept forward through massive numbers of amplifier weights and transfer functions, and the output is the mapping of this *gestalt* into meaningful results. The “intelligence” in artificial neural networks stems from the manner in which the amplifier weights are trained and the specific topology used to map inputs to outputs.

While advances in how the networks are configured and trained have led to fantastic results, there are limits to what can be done if (at runtime) the process is strictly feed-forward. We are all walking proof that sophisticated symbolic processing can arise on a substrate of deeply interconnected neurons whose behavior has rough analogies to the weights-and-transfer functions of artificial neural nets, but it is certain that complicated feedback modes are required for anything remotely in the realm of “higher-level thinking.”

This means that artificial neural networks and other statistical machine-learning technologies are not going to compete for software development jobs soon. Advances in symbolic processing seem to be languishing, except, perhaps, for IBM’s Watson, which seems to be establishing a niche for integrating large-input corpuses to produce correct answers.

Data does not influence the way we develop software. There is virtually no development “best practice” based on statistically meaningful data. And yet, two of the most popular websites in the world are potential treasure troves: Stack Overflow and GitHub. And it’s hard to imagine a task better suited to monitoring for analytics than the computer-based task of writing code.

It’s easy to write scripts that process Stack Over-

flow data by tag or text. While this can be immediately helpful in identifying popularity trends, this is barely scratching the surface, given that every question on Stack Overflow is a testament to something not being clear.

A good example of the way data *should* influence the way we develop software is Microsoft’s “Windows Error Reporting” tool. An article from 2009 by Kirk Glerum and others at Microsoft discusses how these reports are used, and the attempt to properly map a crash report to a single defect “bucket.” When the article was written, this was done using hand-written functions comprising “100,000 lines of code implementing some 500 bucketing heuristics.” This is testimony to the difficulty of the task, but I would be disappointed if, by now, Microsoft Research had not applied machine learning it.

In addition to analytics about the crashes that occur at runtime, what about the errors that are seen at compile time? I recently talked with a colleague about the potential for compiler and IDE vendors to analyze compiler errors. Wouldn’t it be interesting to know which library functions are most frequently called with the wrong parameter types or are correlated with defect-prone functions?

Finally, there’s something fundamentally wasteful about writing unit tests and then writing code to satisfy the constraints you’ve just specified. I like to imagine that a day will come when, for certain programming tasks, I can ask my computer to present me with suggestions. Asking for a function that satisfies defined constraints is a simpler scenario than a generalized software-writing scenario, and I think it’s something we might see within several years.

None of the scenarios I’ve proposed are possible without pre-processing the data into a usable form, as well as some amount of symbolic processing to create the patterns for neural nets to work upon. But software developers have for too long been “the cobbler’s children who go shoeless” and deserve to reap the benefits of this technology. The data sets are there, the deep-learning libraries are open source; all that is required is some insightful developers to put it all together. ■

Asking for a function that satisfies defined constraints is a simpler scenario than generalized software writing.

Read this story on sdtimes.com





John Chang is director of worldwide marketing at CAST.

Guest View

BY JOHN CHANG

Looking behind the UI of mobile apps

The way that large enterprises extend functionality into mobile apps must change the way they think about and manage the development process. Consider for a minute Yelp's mobile app and the one your bank has offered you. On the surface, both apps seem to obsess over ease of use and having visually appealing interfaces. However, just beneath the epidermal UI they are fundamentally different animals.

Banking apps, as with the majority of enterprise mobile apps (expense filing, CRM, ERP or HR apps), navigate through a complex network of Web services and APIs in order to perform seemingly simple operations like displaying an image of your check. They must also deal with the constant threat of security breaches. I'm not saying that Yelp has it easy, but I doubt that Yelp is API'ing with a dozen COBOL programs written in the 1960s just so it can display the address of the closest sushi restaurant.

What's needed is a renewed, fundamental focus on measurement and analysis of mobile applications.

Development today is driven by speed, agility and a "get it out now" mentality. There may be challenges in a given version of an application, so the thinking goes, but they can be rapidly corrected in another revision of the app that's issued tomorrow or next

week. Add to that a laser focus on the customer experience: A recent article in Government Computer News said, "IT shops tend to focus on technical design aspects, but understanding user experience can ensure rapid user adoption." True, but if the application doesn't perform as promised, you can kiss "rapid user adoption" goodbye.

What's missing is an appreciation for the complexity behind the UI and how it affects the technical quality of the app. Is it robust? How do its interactions with internal enterprise Web services affect the stability of the existing enterprise infrastructure? Is it resistant to security breaches? These are questions that IT may be asking, but are too frequently overlooked by the business leaders.

And, as a developer of mobile apps, you may well know the secret that too many others don't: It's far tougher than it looks. Forrester analyst Michael Facemire is quoted in CIO magazine as saying, "People that have done tons of software [work] say mobile is different... It's more common to see folks that say they can do it, [but as the project goes on they real-

ize] it's not really working out." In this environment, it's easy to see how and why incompatibilities erupt. And it's easy to see why trying to pin down the source of a glitch and fix it can be so frustrating.

What's needed, in order to ensure that mobile applications are produced in a way that helps ensure their success, is a renewed, fundamental focus on measurement and analysis of mobile applications and the infrastructure that supports them. Yes, I am talking about going back to the basics, but not at the expense of slowing down the pace of development.

The analyses of mobile applications, properly conducted, can examine the layers of complexity of the different components. For instance, a new mobile application may be required to work seamlessly with legacy user interfaces, or existing middleware in the stack. Too often, different teams manage them... each of them, making their own silo work as well as it can. Injecting new elements, like mobile applications that access services from around the infrastructure, comes with risk to each component.

Even some of the most well-conceived and executed Web services cannot account for the non-functional defects that may be hiding in your legacy infrastructure. At best, the new element may not perform as advertised or desired. At worst, it can create chaos and fleeing users. Neither alternative is desirable.

Again, I'm not saying that the entire development process needs to grind to a virtual halt while each application is manually examined. That's not feasible in today's business environment. The automated analysis of code, however, can rapidly identify potential challenges and vulnerabilities, enabling developers and IT staff to solve them and head off problems before they erupt into full-scale snafus. Visibility across different departments in the IT organization will naturally facilitate collaborative solutions, creating more value in each iteration of the app.

This makes sense from the IT side of the coin, and from the business perspective as well. The cost of application downtime can easily run into the millions of dollars in lost sales at bigger firms; add to that the loss of reputation among customers, who increasingly have their choice of services in any given area. CIO magazine quotes another analyst who says that a badly produced mobile application "ticks off customers. Once they try it and they don't like it, the chance of them coming back is not real high." ■

Read this story on
sdtimes.com



Analyst View

BY ROB ENDERLE



Rob Enderle is a principal analyst at the Enderle Group.

2016: The year smartphones revolted

Here we are at the beginning of another new year and already we are hearing rumblings that the smartphones we have are about to become obsolete. But it'll be different this time, because this year phones are likely to make huge moves due to the massive amount of change in cellular technology.

At the heart of this change is the new Snapdragon Chipset from Qualcomm, the 820, that promises to pretty much make the phone you have look ancient by comparison. We call this kind of change a revolution because it is so massive. Let's go through some of the coming improvements.

Performance: Typically we are lucky to see a 5% performance increase year over year, and the minimum performance improvement people notice is 10%, so when you get a new phone every other year, it seems a tad snappier. With this processor, the hype promises a better than 30% jump in processing performance.

But the even bigger jump is on network performance, which could be more than 200% in some areas, or in places where MU-MIMO routers are installed. That will be a massive bump for things like streaming movies, games, or even your desktop, but just realize both ends have to be updated to see the performance increase. Thus, the 30% is a given; the 200% will depend on where you live and what is installed at work and in your home.

Sound: Where it's phones or tablets, chances are your speakers suck. Even the best, which have two front-facing speakers, tend to have anemic sound because the speaker size just can't handle both high volume and high range. Also, the speakers are so close together that we also get crappy separation.

Well, that changes in 2016 too. With an option to compress the dynamic range, these new phones can pound out the volume and come close to the performance of a desktop speaker able to fill a room (granted, a small room) with music. And using technology similar to what is used to create surround sound, your ears will be fooled into thinking the speakers are much farther apart.

Visual media: I don't see much point in carrying a camera unless you really need the flexibility of a DSLR. Whether you take still pictures or videos, that camera you've had to carry is about to have even less reason to be taken off the shelf. Using the extra power noted above, both movies and still shots

will come out better due to corrective post-processing in the camera. Ghosting, color accuracy and picture clarity, light correction, and picture resolution should improve to where most of us won't be able to tell if the picture was taken by a real camera or a smartphone. The smartphone will also do things that the camera can't, like auto-index the pictures before shipping them to the cloud for safekeeping.

Security: Qualcomm is bringing out an advanced fingerprint sensor that can be put under glass or metal, and because it uses a sonar-like technology, it can see through dirt and still be accurate and reliable. Qualcomm isn't alone, though, as Microsoft is rolling out its Windows Hello feature that scans your eyeball instead of your finger for their phones.

On the new Snapdragon 820-based phone is a behavioral antivirus tool that, because it is hardware-based, should be difficult to overcome and will look for bad behavior. You're still going to need to be careful and avoid side-loading apps, but if something gets through one of the app stores, there is a good chance it'll get caught by this technology.

No more cords: I'm talking wireless everything from power to connecting with your TV or monitor. Granted, initially both will require accessories you don't have now, such as annoying inductive chargers where you have to place the tablet or phone in just the right place. The new technology is called resonance charging, and it can work within inches and be placed underneath desks and tables.

Qualcomm is also promising fast charging that will take your battery to 80% in less than 30 minutes. On top of this, Huawei announced that it had a new battery that could get to 80% in five minutes. That one feature could make their phones really stand out this year.

So blazing-fast phones that are more secure, don't need power bricks or cords of any kind, take far better pictures and videos, and have batteries that can fully charge in minutes. As I look at my own top-of-the-line phone that is only a few days old and soon to be massively obsolete, I'm actually glad this kind of change only happens once in a decade—otherwise my wallet would never survive. ■

Blazing-fast phones that will be more secure, won't need power cords, and take far better pictures and videos.

Read this story on
sdtimes.com





David Rubinstein is
editor-in-chief of SD Times.

Industry Watch

BY DAVID RUBINSTEIN

Seven habits of highly effective DevOps

I recently moderated a webinar in which Sam Guckenheimer, a Microsoft product owner and group product planner, discussed their development journey to DevOps. It was a “lessons learned” presentation that any company looking to benefit from DevOps can use.

To sum it up: The first decade of doing agile development uncovered some key areas. First, your development team needs to be autonomous, yet aligned with the business in terms of goals and quality. Second, managing technical debt is important, as debt interferes with quality and poses a potential risk. Third, monitor the flow of customer

value; that is, make sure you listen to the customer to minimize reworking code. Make sure the stakeholder is satisfied and has input, from idea to production.

These are largely viewed today as essential to agile development. But Guckenheimer said

his developer division found four more “essentials” to moving from agile development into DevOps.

He started with the backlog, saying it is merely a set of hypotheses and beliefs rather than a form of requirements. The hypotheses need to be turned into experiments, which provide the data that can either substantiate or refute the hypotheses. “Learning leads to the next set of hypotheses. It’s not a random walk,” he said.

Guckenheimer shared this story: “We had a signup experience for Visual Studio Team Services. We asked for your identity, we asked you to set up a project, and then we had all these other tiles that informed you about what else you could do. And we were not at all happy with the conversion rate through this screen. So we set up an experiment where we diverted some of the traffic to a new experience. We had a Web version and we have an experience where you sign up through the Visual Studio IDE. In the IDE, we had an instant 7x bump on the new experience for signups from the IDE. On the Web, we had about a 50% bump. But we discovered by the comparison to the IDE that...we were pumping a bunch more traffic to the Web, essentially routing a large volume of unqualified leads in. We wanted to turn off that extra traffic, and when we did, we got to a 3x improvement

on the Web. And of course we went to the new experience based on those results.”

Microsoft could do that, he said, “because we gather evidence in production.” That’s the next essential habit. He said the Microsoft team gathered business-related data (conversions), but also gathered data for troubleshooting, as well as performance metrics, live-site issues that affect end users, and the number of times needed to communicate and mitigate issues. Precise alerting is key to fast issue detection on live sites, he noted, but it’s something people struggle with. You need a production-first mindset.

“We had these Tier 1 people in our service delivery team who were triaging lots of alerts,” said Guckenheimer. “They were looking at 40 or 50 alerts and trying to find which one is the real source of the problem. Based on that, they would contact the responsible individual, the developer on the feature crew who was responsible. He needs to be on the bridge in five minutes if it’s work hours, and 15 minutes if it’s outside of work hours.

“We replaced that with a health model to find...the root cause via machinery. We set up an auto-dialer that would contact the developer responsible. The result is we had a 40x improvement in the month we rolled it out, and reduced escalations by 50%, all by shifting from manual triaging of alerts to putting in place an automation.”

The seventh habit is managing infrastructure as a flexible resource. “We needed to shift the backlog from doing feature work first to doing live site work first,” said Guckenheimer. “We needed a canary instance,” where we work with our closest customers before we toggle the automation to let it move forward to the next rings of deployment.

Microsoft differentiates deployment from exposure with Feature Flags. “All code is deployed, but features not ready for wide usage will have the flag down on them so no users see it, or only select people can see it,” he said. “We have runtime control over how it gets exposed, so we can do experimentation. Customers can lift the Feature Flags on their accounts so they can participate in that experimentation. We can work with dark launches in production and build up experimentation and refinement, but not hide off in a branch and build up debt that we then have to remediate.” ■

Precise alerting is key to fast issue detection on live sites, but it’s something people struggle with.

Read this story on
sdtimes.com





Get to market faster with DevOps

Find your opportunities for growth and see how your progress aligns with Microsoft's seven key habits of DevOps success with a free, 10-minute self-assessment.

Take the free assessment now
<http://aka.ms/devops>

FREE COMMUNITY LICENSE

A \$9,975 VALUE FOR FREE | SUPPORT AND UPDATES INCLUDED

CLAIM YOUR LICENSE

syncfusion.com/SDTimescommunity

Comprehensive offering includes more than 650 components across 13 platforms, an easy-to-use big data platform, and much more.



WEB

ASP.NET
ASP.NET MVC
HTML5/
JavaScript
Silverlight
LightSwitch



MOBILE

Android
HTML5/
JavaScript
iOS
Universal
Windows
Platform
Windows Phone
WinRT
Xamarin



DESKTOP

Windows Forms
WPF
Universal
Windows
Platform



FILE FORMATS

Excel
Word
PDF
PowerPoint



DATA SCIENCE

Big Data Platform
Predictive Analytics



ENTERPRISE SOLUTIONS

Dashboard
Platform
Report Server