GCN

FEBRUARY 2015 VOLUME 34, ISSUE 2

Technology, Tools and Tactics for Public Sector IT

HOW SECURE ARE YOUR OPEN-SOURCE APPLICATIONS?

PAGE 5

BIG DATA'S HEALTH CARE CHECK UP

Legacy datasets, lack of standards, hamper health data analytics

Page 24

NEXT STEPS FOR GIS:

Tools for linking agency geospatial datasets

Page 28

GCN

Technology,
Tools and Tactics
for Public Sector IT

SALES CONTACT INFORMATION

MEDIA CONSULTANTS

Ted Chase
Media Consultant, DC, MD, VA,
OH, Southeast
(703) 944-2188
tchase@1105media.com

Bill Cooper
Media Consultant, Midwest, CA, WA, OR
(650) 961-1760
bcooper@1105media.com

Matt Lally
Media Consultant, Northeast
(973) 600-2749
mlally@1105media.com

Mary Martin
Media Consultant, DC, MD, VA
(703) 222-2977
mmartin@1105media.com

EVENT SPONSORSHIP CONSULTANTS

Alyce Morrison
(703) 645-7873
amorrison@1105media.com

Kharry Wolinsky
(703) 300-8525
kwolinsky@1105media.com

MEDIA KITS

Direct your media kit
requests to Serena Barnes, sbarnes@1105media.com

REPRINTS

For single article reprints (in minimum quantities of 250-500), e-prints, plaques and posters contact:

PARS International
Phone: (212) 221-9595
Email: 1105reprints@parsintl.com
Web: magreprints.com/QuickQuote.asp

LIST RENTALS

This publication's subscriber list, as well as other lists from 1105 Media, Inc., is available for rental. For more information, please contact our list manager, Merit Direct. Phone: (914) 368-1000
Email: 1105media@meritdirect.com
Web: meritdirect.com/1105

SUBSCRIPTIONS

For questions on subscriptions or circulation, contact Annette Levee,
(512) 301-2632 (phone); (512) 301-3361 (fax);
alevee@1105media.com

REACHING THE STAFF

A list of staff e-mail addresses and phone numbers can be found online at GCN.com.

E-mail: To e-mail any member of the staff, please use the following form: *FirstinitialLastname@1105media.com.*

CORPORATE OFFICE

Weekdays 8:30 a.m.–5:30 p.m. PST
Telephone (818) 814-5200; fax (818) 936-0496
9201 Oakdale Avenue, Suite 101
Chatsworth, CA 91311

Editor-In-Chief Paul McCloskey

Executive Editor Susan Miller

Contributing Writers Kathleen Hickey, William Jackson, Stepanie Kanowitz, Carolyn Duffy Marsan, Patrick Marshall, Shawn McCarthy, John Moore, Mark Pomerleau, Brian Robinson

Vice President, Art and Brand Design
Scott Shultz

Creative Director Jeff Langkau

Assistant Art Director Dragutin Cvijanovic

Senior Web Designer Martin Peace

Director, Print Production David Seymour

Print Production Coordinator Lee Alexander

Chief Revenue Officer Dan LaBianca

PUBLIC SECTOR
MEDIA GROUP

Chief Operating Officer and Public Sector Media Group President
Henry Allain

Co-President and Chief Content Officer
Anne A. Armstrong

Chief Revenue Officer
Dan LaBianca

Chief Marketing Officer
Carmel McDonagh

Advertising and Sales
Chief Revenue Officer **Dan LaBianca**
Director of Sales **David Tucker**
Senior Sales Account Executive **Jean Dellarobba**
Media Consultants **Ted Chase, Bill Cooper, Matt Lally, Mary Martin, Mary Keenan**
Event Sponsorships **Alyce Morrison, Kharry Wolinsky**

Art Staff
Vice President, Art and Brand Design **Scott Shultz**
Creative Director **Jeffrey Langkau**
Associate Creative Director **Scott Rovin**
Senior Art Director **Deirdre Hoffman**
Art Director **Joshua Gould**
Art Director **Michele Singh**
Assistant Art Director **Dragutin Cvijanovic**
Senior Graphic Designer **Alan Tao**
Graphic Designer **Erin Horlacher**
Senior Web Designer **Martin Peace**

Print Production Staff
Director, Print Production **David Seymour**
Print Production Coordinator **Lee Alexander**

Online/Digital Media (Technical)
Vice President, Digital Strategy **Becky Nagel**
Senior Site Administrator **Shane Lee**
Site Administrator **Biswarup Bhattacharjee**
Senior Front-End Developer **Rodrigo Munoz**
Junior Front-End Developer **Anya Smolinski**
Executive Producer, New Media **Michael Domingo**
Site Associate **James Bowling**

Lead Services
Vice President, Lead Services **Michele Imgrund**
Senior Director, Audience Development & Data Procurement **Annette Levee**
Director, Custom Assets & Client Services **Mallory Bundy**
Editorial Director **Ed Zintel**
Project Manager, Client Services **Jake Szlenker, Michele Long**
Project Coordinator, Client Services **Olivia Urizar**
Manager, Lead Generation Marketing **Andrew Spangler**
Coordinators, Lead Generation Marketing **Naija Bryant, Jason Pickup, Amber Stephens**

Marketing
Chief Marketing Officer **Carmel McDonagh**
Vice President, Marketing **Emily Jacobs**
Director, Custom Events **Nicole Szabo**
Audience Development Manager **Becky Fenton**
Senior Director, Audience Development & Data Procurement **Annette Levee**
Custom Editorial Director **John Monroe**
Senior Manager, Marketing **Christopher Morales**
Manager, Audience Development **Tracy Kerley**
Senior Coordinator **Casey Stankus**

FederalSoup and Washington Technology
General Manager **Kristi Dougherty**

OTHER PSMG BRANDS

FCW
Editor-in-Chief **Troy K. Schneider**
Executive Editor **John Bicknell**
Managing Editor **Terri J. Huck**
Staff Writers **Colby Hochmuth, Sean Lyngaas, Adam Mazmanian, Mark Rockwell**
Editorial Fellow **Jonathan Lutton**

Defense Systems
Editor-in-Chief **Kevin McCaney**

Washington Technology
Editor-in-Chief **Nick Wakeman**
Senior Staff Writer **Mark Hoover**

Federal Soup
Managing Editors **Phil Piemonte, Sherkiya Wedgeworth**

THE Journal
Editor-in-Chief **Christopher Piehler**

Campus Technology
Executive Editor **Rhea Kelly**

1105 MEDIA INC

Chief Executive Officer
Rajeev Kapur

Chief Operating Officer
Henry Allain

Senior Vice President & Chief Financial Officer
Richard Vitale

Executive Vice President
Michael J. Valenti

Vice President, Information Technology & Application Development
Erik A. Lindgren

Chairman of the Board
Jeffrey S. Klein

INSIDE

FEATURE

CASE STUDIES

BRIEFING

COMMENTARY

GCN (ISSN 0738-4300) is published 11 times a year, monthly except Dec by 1105 Media, Inc., 9201 Oakdale Avenue, Ste. 101, Chatsworth, CA 91311. Periodicals postage paid at Chatsworth, CA 91311-9998, and at additional mailing offices. Complimentary subscriptions are sent to qualifying subscribers. Annual subscription rates payable in U.S. funds for non-qualified subscribers are: U.S. $125.00, International $165.00. **Subscription inquiries, back issue requests, and address changes:** Mail to: *GCN*, P.O. Box 2166, Skokie, IL 60076-7866, call (866) 293-3194, outside U.S. (847) 763-9560; fax (847) 763-9564 or email GCNmag@1105service.com. **POSTMASTER:** Send address changes to *GCN*, P.O. Box 2166, Skokie, IL 60076-7866. Canada Publications Mail Agreement No: 40612608. Return Undeliverable Canadian Addresses to Circulation Dept. or XPO Returns: P.O. Box 201, Richmond Hill, ON L4B 4R5, Canada.

How secure are your open source-based systems?

BY SUSAN MILLER

Responsibility for secure open source software is, well, complicated.

Some believe open source is more secure than proprietary software because, as Linus's Law says, "Given enough eyeballs, all bugs are shallow." That means that the more widely available open software is, the more scrutiny it will receive, the more flaws will be surfaced and the stronger the code will be.

That would be true if components that make up open source code were regularly reviewed and if developers verified the security of components before incorporating them into their work.

But that's not always the case. Like automobile assembly plants that build cars with independently manufactured airbag and brake components, software developers often assume that open source components in their supply chain are reliable, patched and up to date.

Unfortunately, assumptions like that allow for vulnerabilities like those that were behind the Heartbleed bug.

Flaws exist in open source software for a variety of reasons: the components might be old or not mature when they were first used. Or they might not have been audited or adequately tested. But often, once an open source component makes it into a widely used application, it is assumed to be secure, and demand for testing diminishes.

It's not just open source code that's vulnerable. Much proprietary software uses open source components. According to Gartner, 95 percent of all mainstream IT organizations will leverage some element of open source software – directly or indirectly – within their mission-critical IT systems in 2015.

And in an analysis of more than 5,300 enterprise applications uploaded to its platform in the fall of 2014, Veracode, a security firm that runs a cloud-based vulnerability scanning service, found that third-party components introduce an average of 24 known vulnerabilities into each web application.

To address this escalating risk in the software supply chain, industry groups such as The Open Web Application Security Project. PCI Security Standards Council and Financial Services Information Sharing and Analysis Center now require explicit policies and controls to govern the use of components, according to Veracode.

The use of open source in federal systems is also attracting scrutiny. In December, House Committee on Foreign Affairs Chairman Ed Royce (R-Calif.) and Rep. Lynn Jenkins (R-Kan.) introduced the Cyber Supply Chain and Transparency Act of 2014 (H.R. 5793) that would have required any supplier of software to the federal government to identify which third-party and open source components are used and verify that they do not include known vulnerabilities for which a less vulnerable alternative is available.

The bill also would have required the Office of Management and Budget to issue guidance on setting up an inventory of vulnerable software and replacing or repairing known or discovered vulnerabilities. Agencies would have had to annually report on the security of projects using open source components and their suppliers for reference by other agencies.

The bill is important because, as Rep. Royce said in his introductory remarks, much of nation's economy relies on software with open source components.

"It is precisely because of the im-

continued on page 6

Use of open source is attracting the scrutiny of federal legislators, including House Foreign Affairs Committee chairman Ed Royce (R-Calif.) who voiced concerns that the nation relies on the security of the software.

AP IMAGES

continued from page 5

portance of open source components to modern software development that we need to ensure integrity in the open source supply chain, so vulnerabilities are not populated throughout the hundreds of thousands of software applications that use open source components," Royce said.

But not everyone thought the proposed bill was necessary. Trey Hodgkins, senior vice president for public sector at the IT Alliance for Public Sector, told Government Technology that he thought H.R. 5793 duplicated security measures many companies already use.

"We cannot afford to include known exploitable software in our government infrastructure," said Wayne Jackson, CEO of Sonatype Inc., a software supply chain service provider that is the steward of the Central Repository, the largest source of Java components, as well as creator of the Apache Maven project and distributor of the Nexus open source repository manager.

Today, 90 percent of a typical application is composed of open source and third-party components, Jackson wrote in a blog post. The Central Repository clocked in 17.2 billion downloads in 2014 – more than 47 million components every day.

That makes the inventory of open source components critical, Jackson said, because without it, IT managers can't know if their systems contain compromised components.

One way to check is with Application Health Check that provides a free breakdown of every component in an application and alerts IT managers to potential security and licensing problems.

"When open source is found to be defective, it's disclosed, but if you don't know what's in your software, that disclosure tips off adversaries who can use it to exploit vulnerabilities," Jackson said. And hackers get the biggest bang for the buck by going after the components that are widely used, as the OpenSSL/Heartbleed attack demonstrated.

And it's not just enterprise business software that's vulnerable, Jackson said. The problem affects the security of any system with digital components, from websites to cars to insulin pumps. The whole Internet of Things is vulnerable to exploits because it is based, in part, on components that have no upgrade path once deployed.

There may be no way to completely protect government's critical systems from determined adversaries, but ensuring that the basic building blocks are secure is a good place to start. •

CUNY students turn Watson loose on NYC challenges

BY MARK POMERLEAU

Aside from being undisputed world champion of the popular television trivia game show "Jeopardy!," IBM's supercomputer Watson is being used by college students to solve New York City's biggest urban challenges.

Students at the City University of New York (CUNY) competed for cash prizes and IBM internships as they designed apps that use Watson to improve city government services and educational outcomes.

Watson's unique features allow it to solve problems by quickly analyzing huge volumes of data, understanding complex questions posed in natural language and proposing evidence-based answers that help improve decision-making, much like humans.

Using this capability, students built their applications and got hands-on training that will give them valuable technology and business skills necessary to succeed in tomorrow's data-driven workplace.

The first place team in the CUNY-IBM Watson competition designed a virtual case worker assistant to ameliorate hurdles and difficulties facing New York's social workers. The app provides case workers with reports while analyzing various patterns specific to the social work industry. The app's creators believe it will greatly cut down on time spent performing administrative tasks and allow social workers to better serve their constituencies.

The second place team developed an app called SmartCall, which delivers a more organized and efficient 311 information bulletin service to New Yorkers. The SmartCall app will be able to predict complaints using data from the city to resolve concerns faster.

Lastly, the third place team created an education tool called Advyzer that will advise undergraduate students and their counselors and recommend educational tracks based on the student's learning preferences and graduation requirements. The app also takes into account the student's career goals and suggests course schedules based on such information.

"The partnership with IBM offers students the opportunity to look into the future and the way society does business and provides services. It empowers students to shape the future that they will inherit," said Stan Altman, professor at Baruch College School of Public Affairs.

Students competing came from a wide variety of majors such as computer science, marketing, economics, math, urban studies, and finance. All the contestants regardless of where they finished, were able to enroll in summer internships where they can put their applications to use through business and student-led ventures. •

Containers wait on sidelines for government uptake

BY CAROLYN DUFFY MARSAN

While still scarce in government data centers, containerization is a cutting edge technology that eventually promises to make dramatic improvements in the efficiency of cloud-based applications, according to its developers.

Containerization is a standards-based approach to packaging application code that allows software developers to re-use code and create applications that are easily ported to different operating systems and devices.

The approach could make infrastructure-as-a-service more efficient because developers can put more containers on a physical server than is currently possible when running virtual machines.

Containerization is early in its development and is just starting to be used in pilot projects in government. Early adopters include developers of consumer apps such as Yelp, Spotify and eBay. The approach is also being deployed by infrastructure vendors such as RackSpace, which is using the technique for its email service.

Experts believe containers eventually will be deployed in public and private clouds used by public-sector organizations.

"The challenge with containerization is that it's brand new. It's not ready for enterprise, mission-critical apps," said Susie Adams, chief technology officer of Microsoft Federal. "Why Microsoft has invested in it is that it looks really promising. If you can pack more virtual images into a physical server, you can better use your resources. But there is no enterprise-grade management system available yet."

The leading containerization company is Docker, a two-year-old startup that has raised $66 million in venture financing. Docker has gained momentum since it began offering its platform for free under an open-source license, attracting support from Microsoft, Amazon AWS, Google, VMware, IBM and RedHat.

HOW IT WORKS

Docker consists of two parts: Docker Engine, a lightweight portable runtime tool, and Docker Hub, which is a cloud-based service for sharing applications. Docker Engine has been downloaded 100 million times, and 45,000 Docker applications are located in Docker Hub – a sign of the interest in the technology.

Docker applies the concept of the shipping industry's containers to software. A Docker container is a standard way for an application to identify its infrastructure requirements. This approach allows developers to worry only about what is inside the container, while infrastructure operators worry about delivering additional IT resources that the application needs.

Docker containers are designed to be more portable and efficient than virtual machines. While virtual machines consist of an application, binaries, libraries and an operating system, a Docker Engine container includes just the application and its dependencies, running on the host operating system and sharing the kernel with other containers.

"Containerization is really operating system virtualization that is more efficient than the typical approach, and it's more application focused," said Bill Kleyman, national director of strategy and innovation at MTM Technologies, a Stamford, Conn.-based consulting firm. "It provides the necessary resources to run an application as if it is the only app living on the operating system.... App containers have root access, direct access to libraries. This is not something you can do with standard architectures."

One advantage of containers is that they reduce the amount of IT infrastructure – including compute, network and storage – that customers need to purchase in order to run apps.

"Containers are really interesting. It's almost the next version of virtualization that's a little more efficient," Adams said. "Say you want to stand up 10 virtual machines, and each of those virtual machines is 10 gigabytes in size. If you stand up a 10G container in Docker, it shares more resources under the hood. It wouldn't consume even close to 100G of resources because of the way Docker does things."

> "If you can pack more virtual images into a physical server, you can better use your resources."
>
> – SUSIE ADAMS, CTO, MICROSOFT FEDERALL

Containers offer "increased efficiency and increased control over the app you're trying to deliver," Kleyman said. "A containerized app allows for real-time, cloud-native performance. That's the big thing. They offer a greater degree of isolation and looser coupling on layers of virtualization than traditional approaches. This isolation provides a greater degree of reliability and also greater amounts of control."

Government software development shops like GSA's 18F software innovation program are likely to be the first to deploy Docker and other container technologies.

"Any app moving towards web ser-

continued on page 10

GCN

SOFTWARE-DEFINED ENTERPRISE

SLOW BUT STEADY

Moving to a fully software-defined enterprise takes time, and is typically done in steps: virtualization, cloud, software-defined storage, software-defined networking, the software-defined data center, and then full SDE. Slowly but surely, organizations are making their way.

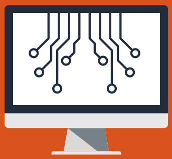

55% of federal agencies are moving to software-defined networking

59% of federal agencies are moving to software-defined storage

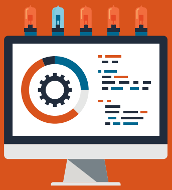

66% of federal agencies are transitioning to software-defined data centers

LOWERING COST THROUGH SDE

One of the biggest pressures on government IT departments is budget. Budgets remain flat at best, yet demands continue to grow. A software-defined enterprise or components of an SDE can help lower IT costs in these ways:

- Software-defined networks can reduce networking costs by **45%** in government data centers
- Software-defined networking can reduce Capex on individual network devices by **50-70%**
- Moving to a software-defined data center can cut Capex by up to **49%**
- Software-defined storage can cut storage costs by **43%** for government agencies
- Software-defined enterprises save organizations **30%** in both operations and development staff time
- SDEs increase revenues by **26%** due to new applications and IT services

CHALLENGES LEADING TO CHANGE

Fast-growing data stores, an explosion in mobile traffic, demand for new services and budget constraints are pressing issues for agencies' IT infrastructures.

Here are the top challenges that a software-defined enterprise can address:

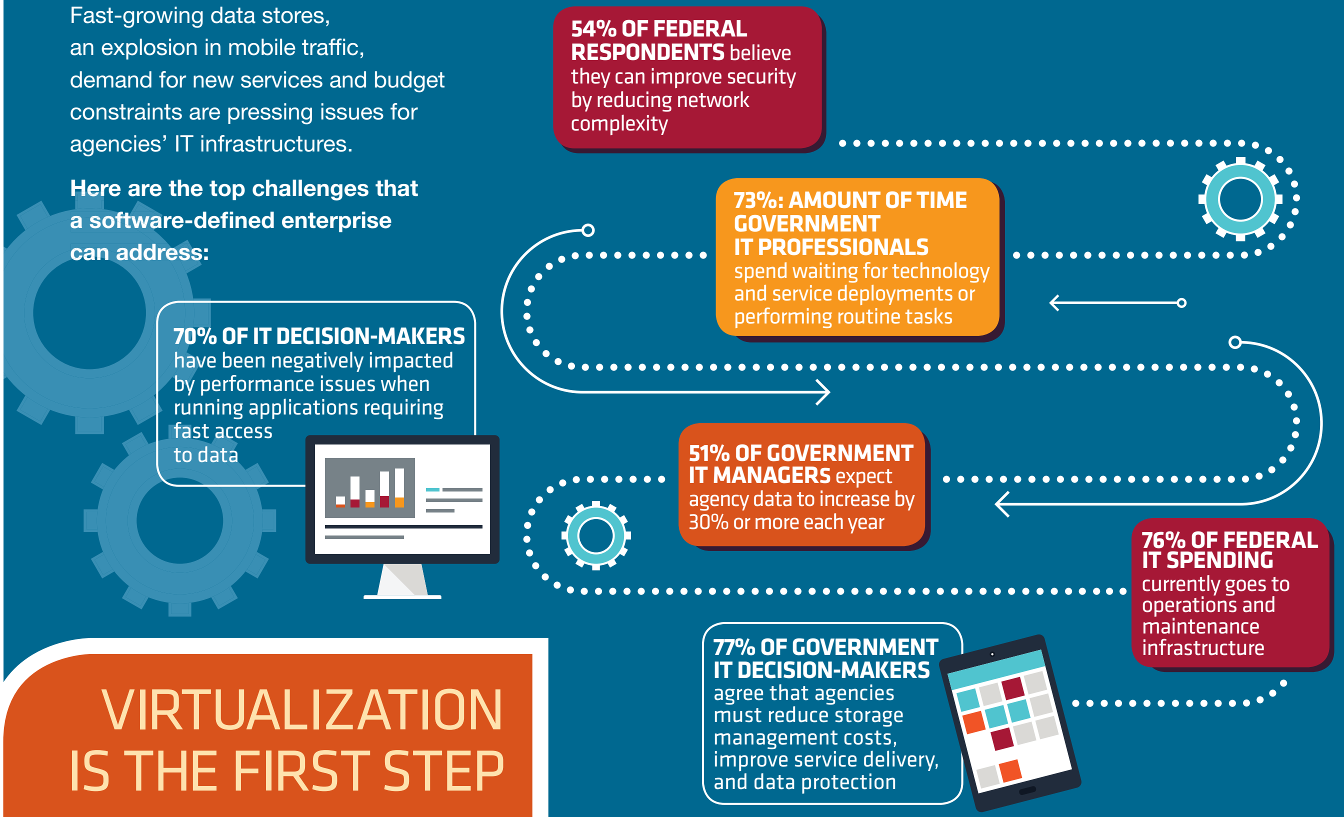

70% OF IT DECISION-MAKERS have been negatively impacted by performance issues when running applications requiring fast access to data

54% OF FEDERAL RESPONDENTS believe they can improve security by reducing network complexity

73%: AMOUNT OF TIME GOVERNMENT IT PROFESSIONALS spend waiting for technology and service deployments or performing routine tasks

51% OF GOVERNMENT IT MANAGERS expect agency data to increase by 30% or more each year

76% OF FEDERAL IT SPENDING currently goes to operations and maintenance infrastructure

77% OF GOVERNMENT IT DECISION-MAKERS agree that agencies must reduce storage management costs, improve service delivery, and data protection

VIRTUALIZATION IS THE FIRST STEP

Federal agencies have implemented some form of virtualization

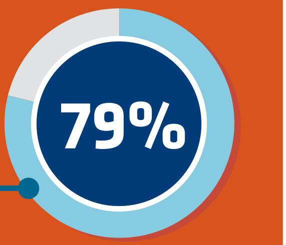

Agencies will have deployed a server virtualization initiative by the end of this year

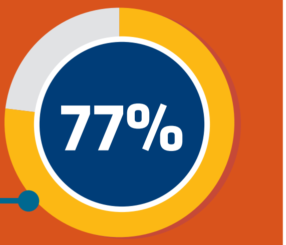

Agencies will have desktop virtualization by the end of this year

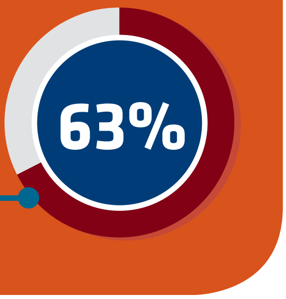

CLOUD COMES BEFORE SDE

Cloud adoption in government is slow but steady, and it's an important step toward SDE. The top three benefits of cloud for government are:

- **BETTER SECURITY**
- **MORE STORAGE**
- **LOWER COST**

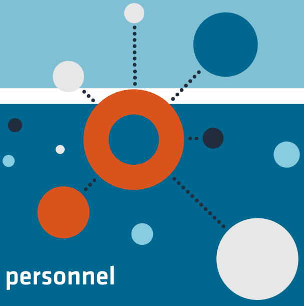

WHY SDE?

- More automation = more effective use of personnel
- Lower cost
- Greater manageability
- Flexibility and agility
- Improved speed and performance
- Faster deployment of new applications

For more information, please visit www.vmware.com/vmwarestory

vmware®

continued from page 7

vices or a Service Oriented Architecture is a good candidate for containerization," said Mark Ryland, chief solutions architect at Amazon Web Services Worldwide Public Sector. "Docker is a standard that has taken off like wildfire. All the major cloud vendors – Google, Microsoft and us – support the Docker container format. The government is going to like it because it is a de facto standard."

Docker rivals include CoreOS, Canonical, Spoonium and Flockport.

The main challenges facing containerization is that the technology is brand new and relatively untested. A good sign that Docker is ready for government apps will be when it is supported in enterprise-class IT management systems like VMWare's V Center, Adams said.

"Right now, you can't spin up the containers or run virtual containers or create high-availability clusters – all the stuff that you want to do to create a virtual computing world," Adams said. "There is no central management console to do it."

In government circles today, Docker is only being used for development and testing applications, but Adams said she has seen government RFPs for development and test environments that request Docker support.

"This is a technology that people are going to want to follow," Adams said. •

Focus on security may cloud awareness of 'shadow IT'

Nearly three-quarters of IT security professionals are unaware of the amount of "shadow IT" within their organizations, according to a recent survey by the Cloud Security Alliance.

Shadow IT, according to CSA, is technology spending and implementation that occurs outside the IT department, including cloud apps adopted by individual employees, teams and business units. "Employees are more empowered than ever before to find and use cloud applications, often with limited or no involvement from the IT department," according to the survey report, which interviewed 212 participants around the world in professional IT security roles.

Some organizations block certain cloud services altogether, such as those from Dropbox, Facebook, Apple iCloud, Tumblr, but that can be even riskier if employees seek out alternatives that have less mature security controls, CSA said.

Shadow IT is not a new problem, nor solely a cloud-based one. For years, removable drives have made it easy to move files from one office to another, but it used to take some ingenuity to get outside the perimeter. When the world went wireless, there was an exponential jump in the ability to think and work outside the box.

But the recent rise of shadow IT might have to do in part with pressure on IT departments to devote more of their time to defending their networks against escalating threats and incursions.

FROM THE READERS

It is not shadow IT, it is clearly what-users-want-at-work IT.

According to CSA, more software vulnerabilities were uncovered in 2014 than any other year on record. And the security of data in the cloud has risen beyond the domain of IT departments and is now a "board-level concern" of 61 percent of the companies surveyed.

IT professionals cited malware as the top security threat facing their organizations (63 percent), advanced persistent threats (53 percent), compromised accounts (43 percent) and insider threats (42 percent).

In fact, cloud security projects were the leading IT project in 2014, according to CSA. Globally, three-quarters of organizations said cloud security projects were very important, moving past intrusion detection and firewalls in the level of seriousness.

The report said that organizations' top concern about shadow IT is the security of corporate data in the cloud, followed by potential compliance violations (25 percent) and the creation of redundant or unplanned services creating inefficiency (8 percent).

Perhaps most alarming, only 8 percent of organizations know the extent of shadow IT at their shops, and 72 percent, "did not know the scope of shadow IT but wanted to know." That number is higher (80 percent) for organizations with more than 5,000 employees, CSA said.

Despite efforts to manage shadow IT, IT departments at 79 percent of firms get requests from their end users each month to buy more cloud applications, according to the CSA survey.

The 2014 Cloud Adoption Practices and Priorities Survey was designed to gauge how IT organizations handled security for cloud services, including how they manage "employee-led" cloud adoption. •

A Cloud Roadmap: From Implementation to Innovation

How a Complete Cloud Offers Agencies a Modern Approach to Cutting-Edge Service Delivery

The government IT infrastructure must evolve. The traditional approach to delivering IT services is proving incapable of keeping up with either the growing complexity of the IT enterprise or the increasing demand for innovation. This pressure to provide an infrastructure that is more manageable, scalable and flexible is pushing government agencies to accelerate their adoption of cloud solutions.

Until recently, government agencies have been slow to move to the cloud. In 2011, the Office of Management and Budget (OMB) issued its "Cloud First" policy, which directed agencies to adapt their IT strategies to take full advantage of the cloud. Yet a September 2014 report from the U.S. Government Accountability Office found that at seven major agencies, investments in the cloud accounted for a mere two percent of IT budgets. Collectively, the agencies consider cloud services for only 33 percent of their IT investments.

However, many government IT executives have not fully understood the benefits of the cloud. In the early days, most discussions about the cloud focused on the potential to lower the total cost of ownership of the IT infrastructure. Such savings are real, as noted by the Government Accountability Office (GAO) report, but they are only part of the story.

In 2014, OMB officials stepped up their efforts to explain the full benefits of the cloud. Innovation emerged as a key selling point. In a traditional IT environment, the process for provisioning server, storage and network systems is cumbersome, making it difficult to field new services in a timely fashion. OMB envisions agencies turning their data centers into "cloud ecosystems," with the ability to deploy infrastructure on demand to support evolving requirements.

Many agencies already have discovered the benefits of the cloud when it comes to supporting their mobile users. Cloud-based solutions make it possible to extend applications and data to mobile devices without compromising on performance. Gartner has predicted that 50 percent of mobile application development will be cloud-based by 2017. Likewise, data-optimized cloud platforms make big data applications more readily available to a broad user base, according to OMB.

Other popular solution areas include employee collaboration, workforce management, and customer or constituent relationship management.

However, to achieve these benefits, government agencies need to develop an enterprise approach to cloud that provides them with a range of solutions that enables them to address existing infrastructure requirements while positioning them to meet unanticipated future needs.

The IT-as-a-Service Paradigm

An enterprise cloud strategy is not just about new technology. It is also about a new mindset.

In its overview of cloud computing, the federal CIO Council highlights the increased flexibility that comes with cloud, including rapid scalability, on-demand self-service, resource pooling, and faster deployment of applications. Taken together, these capabilities make it possible to develop an "as-a-Service" approach to IT, enabling users an agency to tap into a range of IT services on an as-needed basis, scaling up when demand peaks and scaling back when demand recedes.

Cloud service providers offer three basic categories of cloud-based services:

- Software-as-a-Service (SaaS), which offers access to key applications;
- Platform-as-a-Service (PaaS), which provides the underlying IT services for those applications;
- Infrastructure-as-a-Service (IaaS), which delivers compute, storage and related resources.

But such capabilities, in and of themselves, have limited value until agencies adopt an "as-a-service" mindset. That is, agency IT leaders need to develop the policies and processes that encourage and support the development of innovative cloud-based solutions.

Too often, agencies take a piecemeal approach to the cloud, usually beginning with a SaaS initiative, then perhaps expanding to PaaS and/or IaaS, with the different pieces acquired from different vendors and cobbled together as well as possible. In the end, agencies end up delivering a loose affiliation of IT services and asking users to make-do, with everyone paying the price in terms of reduced flexibility, manageability and cost-savings.

In part, these problems have developed because of the piecemeal-like development of the cloud industry. Over the years, a plethora of vendors have popped up offering one service or another, with little thought given to how customers would integrate those services. Perhaps that worked fine in the early days, when most agencies were just piloting individual services, but those days are past.

Today, agencies need to think in terms of an overarching as-a-service strategy. That strategy should incorporate SaaS, PaaS, IaaS and other emerging as-a-service offerings in such a way that users can acquire the services they need, when they need, and never have to worry about how they will work together. Such a framework enables organizations to buy or develop point solutions, but to do so with the enterprise perspective in mind.

Partnering for Success

As might be expected, the development of an as-a-Service strategy has important implications for procurement. If an organization takes a haphazard approach to partnering with vendors—either engaging multiple cloud service providers or allowing individual departments to buy applications independently—the result, simply put, will be a mess.

That's not to say that an organization should tie itself to a single cloud service provider. No one provider, no matter how extensive its portfolio, can offer best-of-breed technology to meet an organization's every requirement. The key is to take a procurement approach that offers both stability and diversity in cloud offerings.

With stability in mind, an organization should look for a cloud service provider who provides the foundational components of an as-a-service strategy—that is, SaaS, PaaS and IaaS—and who has a vested interest in ensuring that those components work together.

"The value of going with a single cloud provider is that it puts that burden on the provider to make sure that your cloud solutions are connected, are talking to each other, and that the integration is not disrupted during the update cycles," said Aaron Erickson, director of Government Innovation at Oracle.

With diversity in mind, the organization should look for a provider who takes an open solutions approach that eases the integration of applications and data from third-party service providers or in-house developers. Stability, in a sense, supports diversity. If application developers have a clear understanding of the underlying cloud services, they can focus their energies not on integration but on innovation.

Finally, public sector organizations need a partner who understands the particularities of their environments. Much of the growth in the cloud industry has been driven by the private sector, where organizations do not have the unique business requirements created by government policies, regulations, and missions. Agencies need a partner who takes those requirements into consideration when developing its solutions.

Oracle recognizes the unique business requirements of public sector organizations and is leveraging decades of industry knowledge and experience in delivering cloud solutions, says Sarah Jackson, vice president, Oracle sales engineering.

"Just like we have done with our other product lines, we invite customers to participate in hands-on validation testing activities and provide opportunities for them to have input into future releases and functional roadmaps," she said.

Oracle Government Cloud: A Complete Solution

Oracle's cloud strategy positions the company as a strong partner for government agencies. The company's global cloud infrastructure includes 17 data centers supporting a comprehensive suite of SaaS, PaaS and IaaS solutions. Today Oracle hosts more than 10,000 cloud customer organizations with more than 25 million daily cloud users. As part of that broader offering, the Oracle Government Cloud provides a series of data centers built specifically to address the rigorous security and compliance requirements of government agencies.

The Oracle Government Cloud provides users access to all of the traditional Oracle technology that agencies are probably already using, including its database and middleware solutions and applications for enterprise resource planning, human capital management, customer management, and project management. It is a solid choice for public sector organizations that must have the ability to deliver modern solutions to citizens and employees efficiently and quickly in order to be successful.

These offerings meet key security and operational requirements common to government agencies, including NIST 800-53 and the Federal Risk and Authorization Management (FedRAMP) program, as well as the International Organization for Standardization (ISO) and the relevant provisions of the International Traffic in Arms Regulation (ITAR).

From a planning perspective, Oracle makes it easier for customers to migrate to the cloud via its Customer 2 Cloud program, which enables Oracle customers to convert on-premise licenses to cloud subscriptions. This program also sends experts into the organization to help evaluate what solutions should move and includes a discussion about costs, benefits and risks to help the customer make the best decision possible.

"Oracle executives recognize that it is important to help agencies feel confident about immediate decisions, understand the steps toward a successful transition and start thinking about possibilities for future services," says Mark Johnson, the company's director of Big Data and Government Cloud.

In the end, the company's holistic approach to cloud technology, as well as its deep understanding of an agency's needs, help set the Oracle Government Cloud apart.

Modern Cloud, Modern Government: Real-World Success Stories

More than 200 public sector institutions in North America and the United Kingdom are taking advantage of cloud services—and Oracle's Government Cloud in particular—to improve services, reduce costs and provide the performance, security and scalability that aren't simply nice to have but required in today's competitive environment. Here are a few examples:

City of Chicago

CHALLENGE

- Review more than 200,000 annual applications for open, city-government positions while adhering to strict hiring regulations and union protocols.

SOLUTION

- Implemented Oracle HCM Cloud Service to better manage a massive volume of resume submissions—as many as 30,000 for a single position—and ensure that the city considers the best qualified candidates.
- Created an objective, auditable process around the development of referral lists and bid lists using the Oracle HCM Cloud.
- Ranked top-tier candidates automatically, based on self-reported qualifications, reducing the number of candidates recruiters must manually screen by an average of 90 percent.
- Reduced average time to fill a position from one year to 90 days and created a cost savings of several million dollars annually, thanks to the Oracle HCM Cloud Service.

Illinois Department of Revenue

CHALLENGE

- Implement a centralized, cloud-based, customer management system to improve customer-service performance and response quality to address ever-increasing inbound e-mail and call volumes.

SOLUTION

- Implemented Oracle Service Cloud to rapidly deliver consistent answers to questions from the public across all channels—enabling taxpayers to get the answers they need to file accurate returns and allowing IDOR to reduce its workload.
- Provided IDOR with a central repository where it can quickly and easily update or modify important tax information when issues arise or policies change, without assistance from internal IT resources.
- Utilized taxpayers' day-to-day questions to help IDOR build more than 500 question-and-answer pairs, enabling the organization to tightly align its online information and the needs of citizens, and making it easy for site visitors to identify specific tax-related topics. Automatically maintained a top-20 question list on the website, ensuring a significant percentage of site visitors can find their answers with a single click.

United States Air Force

CHALLENGE

- Maintain high-quality Air Force personnel services despite mandated resource reductions and replace legacy case- and knowledge-management systems with a more streamlined and efficient solution.

SOLUTION

- Moved many personnel programs to a new web-based, knowledge-management service, based on Oracle Service Cloud technology, helping to streamline operations. Increased monthly use of the tool to approximately 1 million hits.
- Enabled TFSC administrators and call-center agents to access record systems through the agent console, providing access to past cases and enabling agents to attain a holistic view of each customer's complete personnel record, including pertinent dates for evaluations, promotions, and moves.

What is Oracle Government Cloud?

Secure, Innovative Cloud Computing for Government

Built for government agencies, a comprehensive, flexible, and cost-effective suite of cloud applications and technologies.

Why Oracle's Cloud Services?

450 million job candidate records in the HCM Cloud

19 data centers across North America, EMEA and APAC

62 million daily users

23 billion daily transactions

75% of the Fortune 100 runs Oracle's Cloud

Most comprehensive security and compliance standards in the industry including FedRAMP, ISO 27001, HIPAA, ISAE 3402 / SSAE 16, NIST, DIACAP, PCI, ITAR, CFR Part 11

Private, public or hybrid – it's up to you

More than 200 public sector agencies are already in the Oracle Cloud today

Pick One, Pick All

Infrastructure-as-a-Service (IaaS)

such as elastic compute and storage to run any workload in the cloud.

→ **Rapid self-service** provisioning to spin up virtual machines in minutes. Access information and systems from anywhere

Platform-as-a-Service (PaaS)

develop rich government applications.

→ **Database-as-a-Service** — easy to set up, use and manage

→ **Java-as-a-Service** — rapid and agile deployment of any Java application

→ **Document Cloud Service** — secure, web-based Enterprise File Sync-and-Share

→ **Business Intelligence Cloud Service** — load, model, and analyze data quickly and easily

Software-as-a-Service (SaaS) provide enterprise-wide, modern cloud applications to help governments re-imagine their businesses. The best-of-breed SaaS applications in the Oracle Cloud are integrated with social, mobile, and analytic capabilities to help public sector organizations deliver the experiences citizens expect, the talent to succeed, and the agility business demands.

→ **HCM Cloud** — modern HR differentiates organizations with a talent-centric and consumer-based strategy that leverages technology to provide a collaborative, insightful, engaging and mobile HR, employee and executive experience

→ **ERP Cloud** — empower modern finance, procurement and project management with built in Public Sector industry capabilities

→ **Enterprise Performance Management Cloud** — world class planning and reporting with the simplicity of the cloud

→ **Service Cloud** — modern citizen service through unified web, social and contact center experiences

→ **Sales Cloud** — mobile, collaborative, easy and intuitive tools to reach appropriate audiences to deliver on your mission

→ **Marketing Cloud** — personalize every experience to increase efficiency, accuracy and service levels using cross-channel, content, and social marketing solutions with integrated data management and activation

→ **Social Cloud** — enable public sector organizations to provide a better understanding and engagement with citizens and stronger collaboration and efficiencies within the workforce

Mobile dev, code analysis tool debuts for DOD, intell agencies

BY STEPHANIE KANOWITZ

Mobility is still a technology that federal agencies are defining for themselves, with many still playing catch-up when it comes to developing and launching secure applications.

Two companies recently joined forces to help federal agencies meet the need to develop and release secure apps on schedule and within budget. Made available last week, the solution is a joint effort between IT firm CACI and mobile engagement company Appcelerator.

The Appcelerator combines in one tool the Appcelerator Platform's ability to create native, cross-platform apps in JavaScript and CACI's Code Analysis Tool (CAT4) to vet apps for unsafe code that could be exploited.

The customer benefits from having app analytics that can be traced to a customer's investment, said Larry Littleton, CACI's director of mobility solutions, who added that the tool provides the customers with "a specific, metric-driven" return on investment.

The two companies are targeting Defense Department and intelligence organizations because CACI's CAT4 code tool can show compliance with unique federal requirements such as the Defense Information Systems Agency's Mobile Applications Security Guide.

The apps can be customer-facing or mission-oriented to help reduce paperwork for government workers, said Jeff Haynie, Appcelerator's co-founder and chief executive officer. To use the tool, an agency would either contract out app development or handle it in-house using the Appcelerator Platform.

CACI's CAT4, which removes the manual code review process and can work without Appcelerator, can be used in two ways, Littleton said. First, it can serve as a tool for developers, enabling them to perform self-assessed checks during development.

Or, information assurance workers can use "CAT4 reports to review security risk findings as a part of an agency's IA process," he said. When CAT4's analysis finds a problem, "it tells the developer down to the line number how to fix the code," he added.

The Appcelerator Platform includes baseline analytics so that customers can get feedback during the development cycle and after an app is launched with real-time performance, usage and crash metrics. "You get life-cycle analytics from build all the way through deployment and production," Haynie said, providing a complete view into how an app is performing.

This is the first time the firms are offering their capabilities together in one package. "CACI has many mobile app development engagements already using Appcelerator across our market areas," Littleton said. "The partnership expands this relationship to include the insertion of CACI's mobile app vetting tool CACI CAT4."

The move also reflects current market requirements, including getting apps into users' hands quickly, Haynie said. "I think the whole industry, including the private sector, is shifting to much more of an on-demand, get quick ROI from your investments," he said. •

Uber to open its data to Boston transit planners

BY MARK POMERLEAU

The exploding ride-share, taxi-esque company Uber has garnered much praise – and much criticism – for the lack of accountability of its drivers who use their own cars to ferry customers.

Now in an effort to squelch doubters and build partnerships, Uber announced it will partner with the city of Boston to launch what they call a "smart data" initiative.

The plan calls for Uber to release information on its quarterly trip logs, which will include trip time stamps, to assist the city with alleviating traffic congestion and creating smoother transportation systems. Uber will also provide Boston with pick-up and drop-off data on ZIP codes, distance traveled during trips, duration of trips and technical support to interpret and make use of the data.

Uber hopes to establish similar partnerships with other cities across the country. The company maintains that its datasets will benefit city governments because, "most cities have not had access to granular data describing the flows and trends of private traffic. The data provided by Uber will help policymakers and city planners develop a more detailed understanding of where people in the city need to go and how to improve traffic flows and congestion to get them there."

Uber's smart data initiative could be yet another step toward further legitimization with governments, where it is making some headway. Massachusetts recently passed legislation to officially recognize Uber as a mode of transportation, according to a Wall Street Journal report.

However, WSJ also pointed out that New York last year suspended parts of Uber's service for failure to provide data similar to their new "smart data" initiative, which several taxi companies already provide. Uber is also planning to partner with New York City next for its smart data initiative. •

NASA, Microsoft preview holographic computing

NASA and Microsoft broke new ground in scientific computing when they demonstrated software that will give scientists a simulated experience of working together in real time and in three dimensions from the surface of Mars.

Using software called Onsight developed by NASA's Jet Propulsion Lab, and HoloLens, the just-announced wearable, holographic goggles from Microsoft, scientists will be able take images from the Curiosity rover and view them as 3D holograms of the surface of Mars.

The technology "gives our rover scientists the ability to walk around and explore Mars right from their offices," said Dave Lavery, program executive for the Mars Science Laboratory. "It fundamentally changes our perception of Mars, and how we understand the Mars environment surrounding the rover."

Microsoft's HoloLens technology, which the company said would be a part of the its upcoming Windows 10 operating system, is expected to be unveiled later this year and is now being shown in prototype mode.

A demo of the Onsight application by GigaOm's Kif Leswing provided a "very detailed surface replication of Mars, down to the individual rock." Clicking on rocks with an "air tap" gesture enabled the user to explore more of the Mars environment, Leswing said.

According to GigaOm, the device was running several virtual applications, including HoloStudio, a 3D developer application, the Onsight software developed with NASA and a version of Skype.

NASA said the tool will also provide access to the Mars surface and enable engineers to interact with surface features in a more natural way.

The JPL team developing OnSight specializes in systems to control robots and spacecraft. The new holographic software will help them understand the workspace of robotic spacecraft, which can challenging with their traditional tools, according to JPL.

"Previously, our Mars explorers have been stuck on one side of a computer screen. This tool gives them the ability to explore the rover's surroundings much as an Earth geologist would do field work here on our planet," said Jeff Norris, JPL's OnSight project manager.

Using the wearable, holographic goggles, scientists will be able to view 3D holograms of the surface of Mars.

NEWS.MICROSOFT.COM

JPL said the OnSight-HoloLens technology will also be useful for simpler direction of rover operations, which scientists could do by looking at a target and gesturing at different menu commands.

Microsoft said the HoloLens software will support its Windows 10 rollout as a "new generation of Windows" that represents "an era of more personal computing."

At a meeting announcing features of Windows 10, CEO Satya Nadella described the new operating sytem as "the world's first holographic computing platform, complete with a set of APIs that enable developers to create holographic experiences in the real world."

"With Windows 10, holograms are Windows universal apps, making it possible to place three-dimensional holograms in the world around you to communicate, create and explore in a manner that is far more personal and human," he added.

Analysts who were briefed on the technology were bullish on the development.

"If successful, HoloLens will ultimately expand the way people interact with machines, just as the mouse-based interface did in the 1990s and touch interfaces did after the introduction of the iPhone in 2007," said Forrester analyst James McQuivey, in a statement.

JPL plans to begin testing OnSight in Curiosity mission operations later this year. Future applications may include Mars 2020 rover mission operations and other applications in support of NASA's journey to Mars. •

INAUGURAL 2015

ENTERPRISE ARCHITECTURE WEST

SPECIAL GOVERNMENT PRICING ENDS FEBRUARY 24

EA EDUCATION NOW ON THE WEST COAST!

WORKSHOPS: APRIL 20
CONFERENCE: APRIL 21
SACRAMENTO, CA
CITIZEN HOTEL

At Enterprise Architecture West, enterprise architects, project managers and industry experts will convene to discuss contemporary EA and how to apply it to make the mission possible.

Trending topics impacting the EA community to be discussed include:

- Increasing efficiencies and fostering innovation by using cutting-edge EA methodologies
- Using enterprise architecture as a medium for restructuring

KEYNOTE ANNOUNCED!

The Practical Approach to Driving Business Outcomes with EA Today

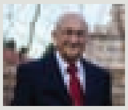

John A. Zachman
Chairman, Zachman International and
Executive Director, FEAC Institute

John A. Zachman is the originator of the "Framework for Enterprise Architecture" (The Zachman Framework™) which has received broad acceptance around the world as an integrative framework, an ontology for descriptive representations for Enterprises.

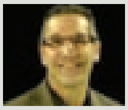

Carl E. Engel
CSO for Zachman International,
President of Zachman Consulting,
and CEO, Elyon Enterprise Strategies

Carl Engel is a visionary who consults, equips and trains today's organizations and their leaders to be tomorrow's success stories.

Don't Miss This Inaugural Event — **Register Before February 24** for Best Savings!

GovEAconference.com
USE PRIORITY CODE: EAW15

PRESENTING SPONSORS

FEAC Institute

ZACHMAN INTERNATIONAL ENTERPRISE ARCHITECTURE

PARTNERING MEDIA

SYSTEMS FCW GCN Washington Technology OMG

PRODUCED BY

PUBLIC SECTOR EVENTS GROUP

Agencies get roadmap for security data sharing

The Office of the Director of National Intelligence's Information Sharing Environment released what it called the first-ever roadmap for national security information sharing, a set of best practices for agencies and IT firms to synchronize data sharing in pursuing national security threats.

The model, called the Data Aggregation Reference Architecture (DARA), was developed over several years as a compendium of ways for agencies to share aggregate information to gain insights into potentially relevant intelligence data, said government executives involved the effort.

"The mission is improving the sharing and safeguarding of information across the whole of government," said Kshemendra Paul, program manager of the Information Sharing Environment, in a recent video statement.

The Information Sharing Environment refers to the people, projects and agencies that enable information national security data sharing.

DARA would provide a reference architecture or model to the groups to blend their information sharing systems and practices. It addresses how to "pull data sets together in a way that protects information security and protects the privacy of individuals who might be represented in that information," said Paul.

The payoffs for agencies and teams using the plan include access to new analytics tools and access to data from other agencies in the national security establishment.

"When analysts are able to search correlated data that other agencies provide using the DARA framework, then organizations do not need to replicate information between systems, which saves storage space, bandwidth and technical staff time across the entire federal enterprise," according to the DARA plan.

The document is also aimed at the national security establishment's industry partners, said Paul, to help them by laying out the roadmap for programs where they will be selling products and services.

Ultimately, DARA is a model for how agencies and firms can link pertinent information across rapidly circulating data streams. Managers of the program said the big data insights produced by the DARA roadmap will save lives.

The Senate Select Committee on Intelligence, for instance, reported that seven of the 14 separate intelligence failures that led to the Christmas 2009 "shoe bomber" incident were "directly related to limits in data interoperability, aggregation and correlation," according to a blog post last month by Michael Kennedy, PM-ISE Executive for Assured Interoperability.

"We recognize that harmonizing the entire federal government and mandating standards cannot happen overnight," Kennedy said, "but the DARA takes an important step in that direction, by establishing consistency in understanding the issues involved, organizational expectations and terminology.

Paul emphasized two key features of DARA, starting with its importance as a basic guide to interoperability. "When you talk about big data analytics, the big dirty secret is the prosaic issues of data plumbing: moving data sets, cleansing data, the basic extraction, movement and loading of information," he said.

"What becomes a core issue is that different agencies, different trading partners, do it differently and that adds a lot of friction, adds a lot of point-to-point connections."

The second take-away from developing the new reference architecture is its maturity model approach.

"Not every agency, not every program is in the same place in terms of their data management approach," said Paul, including "their readiness to share information, to assure policy around security and to show agencies where they are on the continuum of maturity." The roadmap will help them get better in the future, he said.

Paul encouraged agencies to download the DARA and use it to provide feedback. "Let us know where you have had successes and where there's more work we can do to make the guidance that much more useful to you and your program," he said. •

5 goals of the Data Aggregation Reference Architecture

- Define the interoperability requirements for data aggregation enterprise investments.
- Define a reference architecture that enables entity resolution, data correlation and disambiguation across multiple federal databases.
- Provide directions to identify what individual agencies need to do to embrace a federated approach and assess the possible organizational impacts.
- Specify what individual agencies need to do to embrace a federated approach and possible enhancements to their investments.
- Serve as a broad, general reference architecture that guides the creation of more specific, concrete solution architectures.

Alliance proposes low-power wireless protocol for IoT

A group of computer and networking firms have formed an alliance to standardize on the use of low-power wide-area networks, or LAPWANs, to drive development of Internet of Things (IoT), machine-to-machine and smart city applications.

The LoRa Alliance is dedicated to using protocols derived from LAPWAN, including LoRaWAN, to ensure interoperability of IoT applications between telecom operators and other firms who have joined the effort.

The LoRaWAN technology is considered suitable for IoT and M2M applications because it extends much farther than cellular technology and often operates in small sensor-type devices that can last for months on the power of a small battery.

LoRaWAN also lets public networks link multiple applications using the same network infrastructure, which will help enable new applications for IoT, M2M, sensor networks and industrial automation applications, according to the group.

Device manufacturers and developers are also using the technology to propose solutions at a lower total cost of ownership and with longer battery life that often do not need a powerful cellular connection, according the Alliance's announcement.

"The LoRa technology is ideal to target battery-operated sensors and low-power applications as a complement to M2M cellular connectivity," said Richard Viel, chief operating officer of Bouygues Telecom, a French mobile phone service provider.

> "With LoRaWAN, entire cities or countries can be covered with a few base stations."
>
> – OLIVIER HERSENT, ACTILITY

Olivier Hersent, CEO of Actility, a French energy management and IoT services firm, said the low power technology protocol enables M2M applications to scale across long distances on low cost networks.

"With LoRaWAN, entire cities or countries can be covered with a few base stations, no longer requiring the upfront rollout and maintenance of thousands of nodes as in traditional mesh networking," he said.

"This has made IoT possible now, with minimal infrastructure investment."

Open technologies are also a key enabler of M2M connectivity, according to Thorsten Kramp, master inventor at IBM Research, another member of the alliance.

To that end, IBM has released the IBM protocol – LoRaWAN in C – as open source, he said, "which provides a solid foundation for the development of a broad range of end devices compliant with the LoRaWAN specifications."

"To encourage the mass adoption of low cost, long range, machine-to-machine connectivity, open ecosystems are critical," Kramp said.

Gartner analyst Nick Jones told IDG News Service's Stephen Lawson that the IoT technology market is "in what you might call a land-grab phase. Everyone is trying to get ahead to establish their presence."

Prospective alliance members include Actility, Cisco, Eolane, IBM, Kerlink, IMST, MultiTech, Sagemcom, Semtech, and Microchip Technology. Telecom operators so far include Bouygues Telecom, KPN, SingTel, Proximus, Swisscom, and FastNet, from South Africa. •

City smoothes disaster recovery upgrade

Riverside, Calif., wound up a project to upgrade its disaster recovery infrastructure that resulted in doubling its data storage capacity and enabling fail-over capabilities, all without disrupting data access for its end users.

The data storage and security project involved migrating half a petabyte of data from the city's outgoing storage area network. A NetApp FAS8040 storage array was deployed at the city's primary data center, and a NetApp FAS8020 array was installed at another location for disaster recovery.

NetApp SnapMirror replication software was used to provide a duplicate copy of the city's data at the disaster recovery site, while the firm's Clustered Data OnTAP storage operating system enabled failover and scalability.

The city contracted for the design of the system, deployment and data migration services with DataLink Inc., a provider of data center infrastructure services.

The new system provides nearly a petabyte of storage that can be scaled up and has the ability to move workloads to any storage pool within the cluster to accommodate changing requirements or hardware refreshes, according to Datalink. The features were put in place entirely after hours to avoid service interruptions.

Datalink is also providing technical support through its OneCall service, which is staffed by two U.S.-based support centers to ensure redundancy and compliance with service-level agreements. •

Low-tech sectors to see more IT spending

BY MARK POMERLEAU

Cash-strapped state, local and education (SLED) agencies started feeling the budget pinch around April 2014 and began reeling in their IT spending compared to the previous year.

But while IT departments were decreasing their investments, other areas like education, law enforcement and road construction have been "using technology to better meet their objectives while reducing overall (non-IT) costs," according to a recent report by Onvia, a government business development consultant.

More communities are investing in body cameras to document the behaviors of public safety officers as a means to increase accountability. Reports indicate body camera technology has doubled from 2013 to 2014. In fact, President Obama recently requested $263 million for body cameras at the state and local level. In the past, municipalities have paid between $50,000 and $1 million for body camera contracts, and there are potentially 9,000 departments that are interested in similar procurements in 2015, Onvia said.

Similarly, communities are also projected to increase investment and procurement of school bus cameras to ensure greater student safety. Most commonly, buses are outfitted with three cameras – inside and outside – and some communities have invested in equipping their entire bus fleet with them.

The education sector has rapidly increased the use of tablets and laptops to foster technology-based learning in schools. Tablet contracting increased 21 percent between 2013 and 2014, and this growth is expected to continue. In 2013, 85 percent of Chromebooks sold were placed in school systems, which numbered 2.5 million devices. Other vendors, like Curricula, have focused on bringing technology such as 3D printing into the classroom.

As more people primarily use their mobile devices to access the Internet, governments are making their services and websites more mobile friendly. As such, state and local governments are investing in open data and engagement tools as well as crowdsourcing technology to help drive innovation.

Procurement of intelligent transportation systems has increased by 13 percent, Onvia reported, a trend that is also expected to continue. Intelligent transportation systems are used by state and local governments to alleviate traffic congestion through a combination of sensors, computers and fiber optic networks that update traffic signals in real time based on the current traffic. •

Harrisburg U builds cybersecurity center for state

The Harrisburg University of Science and Technology's Government Technology Institute has established a new center focused exclusively on safeguarding government data and systems from unauthorized access.

The Security Center of Excellence (SCoE) is believed to be the first such center focused solely on securing data entrusted to state, county and local governments, the university said.

Cisco, Deloitte Consulting, IBM, Symantec and Unisys have all agreed to sponsor the SCoE and bring their global experts to HU to help the institute showcase the benefits of collaboration among cybersecurity experts from government, academia and the private sector.

"Our goal is to make this a national best practice for training and supporting those within government responsible for safeguarding sensitive data," said Barb Shelton and Charlie Gerhards, co-directors of the institute.

> "These are some of the best security companies in the world, and they will clearly help this Center to achieve its goal."
>
> – ERIC DARR, PRESIDENT OF HU

Eric Darr, President of HU, said "these are some of the best security companies in the world and, they will clearly help this Center to achieve its goal and in turn help Pennsylvania's government safeguard citizen data."

"It also is a tremendous opportunity for our faculty and students to work closely with government IT leaders and distinguished experts from the technology companies that have agreed to help Pennsylvania continually improve cybersecurity," he added.

The educational program for security specialists in government is planned to begin in spring 2015 and will be followed with seminars, technology testing and collaboration among multiple levels of government.

Chuck Davis, a corporate faculty member at Harrisburg and a security expert, noted that "the number of cyberthreats to companies and governments is increasing exponentially and we constantly need to do more to protect confidential data." •

Visual Studio® LIVE!
EXPERT SOLUTIONS FOR .NET DEVELOPERS

vslive.com/lasvegas

Las Vegas MARCH 16 – 20
BALLY'S HOTEL & CASINO LAS VEGAS, NV

LAS VEGAS
Code Trip
NAVIGATE THE .NET HIGHWAY

Code on the Strip

Visual Studio Live!'s first stop on its 2015 Code Trip is Las Vegas, situated fittingly near Historic Route 66. Developers, software architects, engineers, and designers will cruise onto the Strip for five days of unbiased and cutting-edge education on the Microsoft Platform. Navigate the .NET Highway with industry experts and Microsoft insiders in 60+ sessions and fun networking events – all designed to make you better at your job.

TRACKS INCLUDE:

- Visual Studio / .NET
- JavaScript/HTML5
- ASP.NET
- Cross-Platform Mobile Development
- Windows 8.1/WinRT
- Database and Analytics
- Cloud Computing
- Windows Phone

Register NOW and Save $500!
Use promo code VSLDEC1

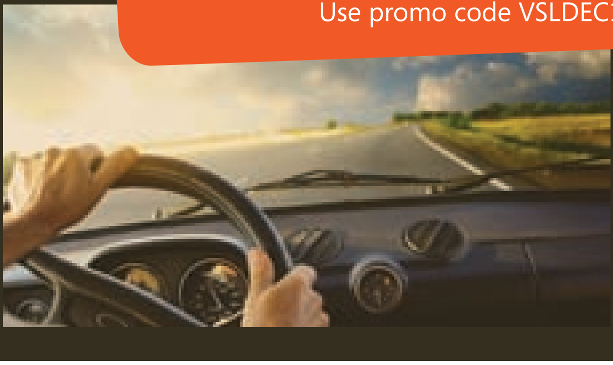

Scan the QR code to register or for more event details.

SUPPORTED BY
Microsoft Visual Studio msdn magazine Visual Studio MAGAZINE

PRODUCED BY
1105 MEDIA INC

vslive.com/lasvegas

BY WILLIAM JACKSON

Police body cameras are only one piece of the video equation

PRESIDENT OBAMA RECENTLY proposed a $263 million program for training and equipment to help make police departments more accountable after the high-profile incidents of police violence in Ferguson, Mo., New York City and elsewhere. It includes $75 million over three years to help purchase 50,000 wearable body cameras for officers.

The cameras are lightweight, high-resolution alternatives to dashboard-mounted video systems already in use in many police cruisers. But the new cameras are only the first step in supporting a new video system for police. Once the cameras have been bought, departments will have to store, manage and secure terabytes of data, sometimes for decades. And because the quality of the video the cameras produce is improving, the amount of storage needed is likely to be much greater.

Many departments are familiar with the requirements of older black and white surveillance and dash-mounted video systems. But new body cameras can produce high-definition, full-motion, wide-area color images. This is great for evidence but can quickly overwhelm current storage systems.

One camera can produce 2.3 gigabytes per hour, or 18.4 gigabytes per shift, said Dave Frederick, senior director of product marketing for Quantum, a storage solution provider. Because most cameras will be activated only during an incident, they probably will not be used eight hours at a stretch. But they still could produce 9 gigabytes of data per shift, he said.

How much storage this will require depends on a number of variables: The number of cameras in use, their format and resolution and the video retention policies. Department policies can call for saving video for anywhere from a month to a year, but if the video is used in court, rules of evidence can require that it be kept for years and – in the case of a conviction – possibly for the lifetime of the defendant.

One department with 1,500 officers found that it would need 700 terabytes of storage to accommodate body cameras, Frederick said, more than double what was needed by its older dash-cam system.

Quantum proposes a tiered-storage solution that balances the cost of storage with performance. Such a system typically would include a high-performance ingest system, using more expensive spinning disks to quickly take in new video and make it accessible.

Subsequent storage tiers could include lower-performance spinning disks and tape for archival storage, which is not as fast but is less expensive. Making such a system practical depends on the ability to automatically move data from one tier to another as needed, without violating chain of possession rules for video used as evidence.

Another consideration is that video being archived as evidence might have to be readable 25 years or more from now, when technology is likely to have changed dramatically. Anyone who has been stuck with a shelf of Betamax tapes can appreciate the challenge of future-proofing video archives.

The cloud could be an attractive short-term solution for storing police video, as long as security and management requirements can be met. But at some point, the long-term cost of renting storage space is likely to overwhelm the initial savings, and any department expecting to use video for the long haul will have to decide how best to acquire and manage its own storage system.

These challenges do not mean that police departments cannot or should not take advantage of new video technology to better document activities on the street. But they should remember that the camera is only the front end of a larger system, and once the "record" button is pushed, there will be an obligation to manage the video for years to come. •

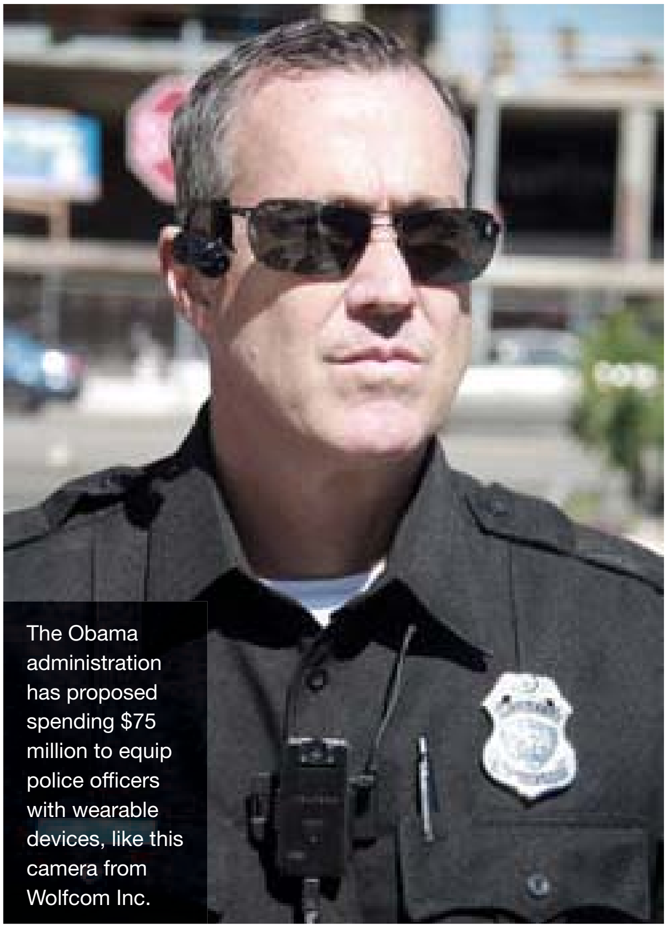

The Obama administration has proposed spending $75 million to equip police officers with wearable devices, like this camera from Wolfcom Inc.

WOLFCOMUSA.COM

INDUSTRY INSIGHT

BY JOSE CARLOS LINARES

Are you ready for the next OS migration?

MORE THAN NINE MONTHS have passed since Microsoft terminated all support for Windows XP. Even though many government agencies successfully made the transition to Windows 7, it's already time for IT managers to start planning for the next operating system migration, as mainstream support for Windows 7 just ended on Jan. 13.

Even though Microsoft will continue to provide security patches until Jan. 14, 2020, there's no time like the present to start planning ahead to protect your infrastructure.

> Even though many government agencies made the transition to Windows 7, it's already time to start planning for the next migration.

Many agencies delay migration because it is a huge task and it can be risky, time consuming and expensive. The time required for a migration depends on agency size and the number and types of applications. It is recommended that IT departments start planning at least one year in advance. A complete migration will cost an agency at least $250 per machine.

Fortunately, government agencies can save both time and money while continuing to provide essential services to citizens by using automated tools and a four-step migration process:

Step 1: Planning ahead: Inventory, analysis and rationalization. The migration process cannot begin until all applications, hardware and network users and resources within an organization have been verified. IT managers can use automated tools to take inventory and compare application files to an established compatibility database.

Incompatibility between applications, hardware and the new operating system can be identified, and the content to be migrated can be rationalized. The time to completion can also be estimated based on the amount of data to be migrated. These programs also help to establish and maintain an inventory of all systems and software, while providing valuable information about the users of each application.

Step 2: Applications: Testing, remediation and repackaging. Once IT managers have determined which hardware and software will be migrated, they must make sure that applications work with the new operating system. Automated compatibility testing helps to identify areas of risk early in the process and allows for remediation and repackaging. Customized reports also help IT staff easily find and fix compatibility issues in a matter of minutes, instead of days, weeks or months.

When deciding whether to migrate legacy systems, IT managers must weigh the value of the application and its importance to the agency's mission. Perhaps only a few employees use the application, or maybe the person who developed the program is no longer available to support it. Agencies must also make sure the browser is certified as secure and operational with the new operating system. Citizen-facing applications may require a different browser than an internal facing browser.

Step 3: Deployment. To successfully move an end-user system to a new operating system, several time-consuming processes must be completed. It is estimated it takes most organizations at least an hour to reimage and deploy a single computer with a new operating system. To help with the process, a systems imaging solution can be used to create several "master" images that can be used on multiple machines.

Additional updates and personalization can be added on top of the image as required. A centralized deployment system can install images on remote computers where on-site IT support is not available. Other automated systems can complete complex tasks such as unattended deployment after hours.

Additionally, a cloud solution that allows end users to self-manage the migration by scheduling a convenient appointment is a better use of everyone's time. That way, IT staff can actually spend time providing essential support, and end users can plan ahead for the interruption.

Step 4: Support: Post-migration maintenance. Once a user's system is migrated, automated tools can support the new operating system environment. Through this technology, systems can be tracked, updated, secured and managed on a consistent basis. This helps IT managers meet compliance requirements and make sure that the new environment will continue to provide efficiencies.

Even though 2020 seems like a long time from now, it's only five years away and will certainly sneak up fast. By planning ahead and using the right tools and processes, IT managers can use a migration to increase agency efficiency while saving time and money. •

— Jose Carlos Linares is CEO of the Open Technology Group.

BY (ISC)2 GOVERNMENT ADVISORY BOARD EXECUTIVE WRITERS BUREAU, LOU MAGNOTTI

Is antivirus software still relevant?

FEDERAL AGENCIES are big users of antivirus software, and regardless of their technical competence, government security professionals still find themselves victims of malware. Unfortunately, simply installing antivirus technology does not protect today's endpoints.

In a 2014 Lastline Labs study on the effectiveness of antivirus scanners, much of the newly introduced malware went undetected by nearly half of the antivirus vendors. After two months, one third of the antivirus scanners still failed to detect many of the malware samples. The malware dubbed "least likely to be detected" went undetected by the majority of antivirus scanners for months or was never detected at all.

For those malware samples that initially eluded all of the scanners, the average time for at least one of the samples to achieve detection was two days. None of the antivirus caught every new malware sample.

No matter how useful antivirus software can be, its drawbacks are causing information security professionals to take a second look at antivirus protection – and the alternatives.

Several years ago, the Milnsbridge Corp. sponsored case studies focused on a new approach, called CloudAV that moves antivirus functionality into the network cloud and off personal computers. The study focused on virtualizing the detection functionality with multiple antivirus engines, significantly increasing overall protection.

Traditional antivirus software that resides on most PCs checks documents and programs as they are accessed. Because of performance constraints and program incompatibilities, only one antivirus detector is typically used at a time. CloudAV, however, can support a large number of malicious software detectors that act in parallel to analyze a single incoming file. Each detector operates in its own virtual machine, so the technical incompatibilities and security issues are resolved.

No matter how useful antivirus software can be, its drawbacks are causing security professionals to take a second look at the alternatives.

Some of the drawbacks deal with speed in handling the volume of data. While CloudAV stores previously screened data, processing time is an issue. There is also the concern of the cloud provider's level of security in and of itself. Regardless, several CloudAV providers are available in today's market.

Many of the existing operating systems come with antivirus software built in. Others may use application whitelistings (AWL) – as opposed to blacklisting – as an integral part of the OS.

Most people in the IT field are familiar with blacklisting because it is the technology used in almost every antivirus product in existence. It simply checks every new file on a system to see if it contains malware. If malware is detected, it is blocked from executing and carrying out any damage.

AWL is just the opposite. It will deny the execution of any application not previously and explicitly identified as "not malicious." AWL offers more security primarily because it denies malicious code that has never been seen before (zero-day issues) and code that blacklists won't recognize immediately. Security professionals must keep in mind that there is considerable expense in the AWL game, not only with the initial purchase but with the internal man-hours required to make changes and test new patches and application updates on the servers.

Another reason information security professionals are taking a second look at antivirus protection is the "cost vs. rewards" to their respective organizations. The advent of malware insurance has offset the cost incurred by damages from malware; however, there are also losses to one's reputation and possibly even regulatory fines to consider. Couple this with the premise that no antivirus technology will guarantee 100 percent security, and government security professionals find themselves in a conundrum when faced with the task of providing cost-effective advice to senior executives.

So, what is an agency to do? While the drawbacks of using antivirus are all valid, many agree that the technology should still be used as part of a "security-in-depth" approach. Maintaining an arsenal of sophisticated security tools that protect the enterprise network from the "outside-inward" is still the preferred, balanced approach to security. Equally important, antivirus technology must be complemented with a good security education and awareness program along with other information security policies and procedures. •

— Members of the (ISC)2 U.S. Government Advisory Board Executive Writers Bureau include federal IT security experts from government and industry.

Lou Magnotti, Executive Writers Bureau member, was lead author of this peer-reviewed article.

The Changing Cyber Threat Landscape

In less than a decade, more than 87 million records with sensitive or private information have been exposed due to cyber-incidents on federal networks alone . But the sheer volume of cyber-attacks is just part of the challenge facing federal agencies when it comes to network security.

Even more worrisome is the ingenuity of today's cybercriminals. Whether it's an insider threat or an external force wishing to do harm, the techniques of cybercriminals have become more insidious and intelligent. Once a threat is mitigated, 10 more will rise up in its place. Take, for example, the increase in backdoor techniques attackers are using to steal sensitive data. These methods let attackers enter a network via an unmonitored loophole and bypass firewalls and anti-malware through sophisticated methods. They can check for available connections and transfer files, connect via social media sites, use custom DNS lookup to bypass detection and even change protocols.

INCREASINGLY, ORGANIZATIONS ARE BEGINNING TO REALIZE THAT PERIMETER-BASED SECURITY JUST ISN'T ENOUGH ANYMORE.

In addition to more aggressive, innovative methods that enable attackers to bypass the perimeter, the very nature of what is inside the network and what is outside the network has changed. For example, remote employees or teleworkers must enter the network to perform their tasks. Are they considered to be inside or outside the network? The same questions are true about mobile users, who use wireless technologies to access the network, contractors and partners, and cloud computing solutions.

All of these types of users and solutions require defining the perimeter in a new way and, as a result, also require reconsidering network security methods. Another related issue is that many agency networks have older devices, technologies and applications that simply aren't able to manage today's cyber-related challenges. They sometimes can't, for example, distinguish appropriate traffic or credentials from malicious traffic or false credentials; keep track of all encrypted traffic, especially in virtualized infrastructures; or adequately filter network traffic.

Increasingly, organizations are beginning to realize that perimeter-based security just isn't enough anymore. Keeping an agency's data and applications safe today requires re-evaluating and updating network security plan, technologies and policies to reflect today's changing threat landscape.

CYBERSECURITY BY THE NUMBERS

1.5	The number, in millions, of monitored cyberattacks in the United States in 2013
15	The percentage by which the average cost of a data breach increased in 2014
21	The percentage of federal breaches in 2013 caused by government employees who violated policies
24	The percentage of government agencies reporting that operating systems or files had been altered, and/or unauthorized access or use of data, systems and networks
25	The percentage of security professionals who doubt that their organization has invested enough in cyberthreat defenses
37	The percentage of cyber incidents that aren't detected by civilian agencies
42	The percentage by which the number of reported federal network breaches increased between 2009 and 2013
49	Percentage of organizations that don't perform employee security awareness training
260	The number of days, on average, that it takes organizations to detect and respond to insider attacks, compared to 170 days for other attacks
46,605	The number of breaches of federal computer networks in 2013

Network Security in the Age of the Software-Defined Data Center

Over time, more agencies are embracing the software-defined data center, citing benefits like lower costs, efficiency, scalability and reduced downtime. With a software-defined data center (SDDC), all infrastructure—servers, storage, and networking—is controlled by software and is automated, virtualized and delivered as a service. Moving to a software-defined data center also allows agencies to more easily embrace trends and technologies like cloud computing, virtualization and converged infrastructure.

The software-defined data center also provides several security advantages. Virtualization means that security services are more pervasive, and automation of the SDDC means that security functionality can be embedded into virtual switches. Policy enforcement is also more effective, since security policies are automatically attached to specific workloads, even as they scale up or out. And in an SDDC, network control is centralized and automated, providing more visibility into network behavior.

In the most advanced software-defined data centers—those that employ network micro-segmentation—administrators can manage networks much more granularly. Micro-segmentation enables unit-level security controls to be implemented in a scalable and cost-effective manner both within and between data centers.

SECURITY AND THE SOFTWARE-DEFINED NETWORK

Software-Defined Networking—a subset of the software-defined data center that is often the first step toward a full SDDC—is also a major boon for data center security.

A software-defined network (SDN) separates the controller, which routes packets, from the data plane, which forwards network traffic to its destination. It is centrally managed via software, enabling more nimble traffic control. The level of automation in a software-defined network eliminates much of the manual configuration work humans generally do—and along with it, the inevitable mistakes that people can make. The automation of policy enforcement also helps ensure that no security protocols are inadvertently breached.

Because the network is fully controlled by software, network administrators can configure the network such that all traffic—both perimeter and internal traffic—is routed through one firewall and intrusion prevention system. If a threat is identified, the software can quarantine affected areas of the network and reconfigure the network in whatever way is necessary to mitigate the problem.

NEW THREATS REQUIRE NEW APPROACHES TO SECURITY

FOR DECADES, THE SECURITY MANTRA has always been “Trust, but Verify.” That’s just not good enough anymore. The new reality in cyber-security is “Never Trust, Always Verify”. That means that no person, device, packet or application, either inside or outside the network, can be assumed to be trustworthy.

The reason for this new, no-holds-barred approach is clear: traditional security methods, which focus on a network’s perimeter and consider anyone inside the network to be trustworthy, have let us down.

A new approach to network security, called “Zero Trust”, assumes that no network traffic can be trusted. The goal of the model, developed by Forrester Research, is to ensure that all resources are accessed securely regardless of location, minimize allowed access to resources as a way to reduce the pathways available for malware, and inspect and log all network traffic.

The Zero Trust model focuses on securing the data first and the network second. That means that security travels with the data, both inside and outside the network. The data-centric approach ensures that even when the data leaves the environment because it is accessed by a mobile device or shared with a partner or contractor, it remains secure. The Zero-Trust approach achieves data-centric security by requiring that files, emails and other data is classified automatically and continuously. That way, if the value of data changes—it becomes more sensitive, for example—the data will be automatically reclassified to reflect that status.

The Zero Trust security model requires securing access to all network resources in a different way from the traditional approach of firewalls, intrusion prevention, content filtering and encryption. The most important change is the requirement for network segmentation, which enforces the separation of traffic. The Zero Trust model recommends using segmentation gateways, which segment networks based on the type of data traveling through it.

2015: The Year of Microsegmentation

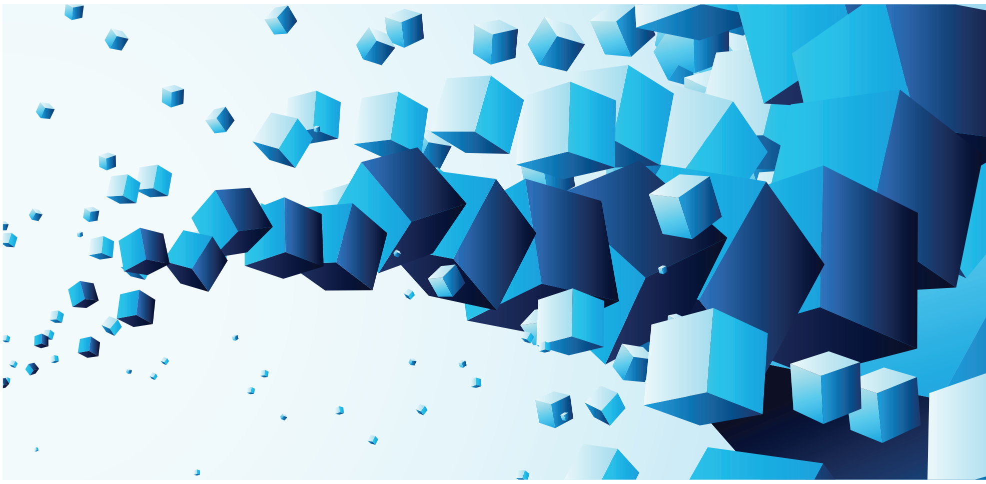

Keeping government data and networks secure is a constant challenge for agencies, made more complicated by ingenious hackers and constantly evolving threats. All levels of government are working hard to keep up. At the federal level, agencies including the Defense Department are asking industry how to better identify and prevent cyberattacks, the National Institute for Standards and Technology (NIST) continues to improve its Cybersecurity Framework, and the House of Representatives recently added a new subcommittee to focus on cybersecurity. At the state level, the top priority for state CIOs in 2015 is security—the same as it was in 2014.

As agencies look for innovative solutions to pressing cybersecurity issues, they are beginning to understand that status quo is no longer good enough. While government has a strong security defense at the network's perimeter, new threats can still get through, often by latching on to legitimate user access. And once inside the network, hackers can jump from server to server, causing untold harm.

"We have seen several incidents the last few years that clearly indicate that the traditional approach to cybersecurity isn't working," said Ahmed Ali, Networking and Security Architect at VMware. "Agencies need a better way to harden their cybersecurity posture, achieve secure multi-tenancy for shared service data centers, and securely take advantage of cloud computing."

MOVING AWAY FROM A PERIMETER-ONLY DEFENSE

Limiting unauthorized lateral movement within a network—one of the biggest cybersecurity challenges today—is one of the most difficult challenges for most data centers. That's because most firewalls control the network via physical points. When traffic passes through these control points, the firewall either blocks the traffic or allows it to pass through. While this method works, it is limited in what it can accomplish. Not only does it require continuously adding more physical firewalls as data center requirements increase, but traditional firewall technology requires manual modification of firewall rules each time a new virtual machine is added, moved or decommissioned.

For many data centers, the solution is microsegmentation—the concept of segmenting network traffic even more granularly to catch intra-network security breaches. In addition to segmenting north-south network traffic (such as an intruder trying to infiltrate a network), it also segments and protects east-west traffic (the lateral traffic that occurs inside the network). This approach is fully in line with Zero-Trust Security; it trusts nothing on the surface.

Network virtualization is a microsegmentation platform with advanced security features that works with external firewalls, such as those from Palo Alto Networks. Ideal for software-defined data centers, network virtualization enables security managers to segment network traffic down to the virtual Network Interface Card (vNIC) and monitor inter-NIC traffic.

With this approach, the firewall is integrated into the network virtualization platform and operates in the kernel of every hypervisor. The platform is fully automated, from provisioning to workload and policy changes. Enforcement is distributed at every virtual interface and within each kernel. Network virtualization delivers about 20 Gbps per host. That means that every time the data center adds additional hosts or workloads, the platform adds additional firewall capacity for the east-west traffic.

The combination of traditional firewalls for north-south traffic and a microsegmentation solution like VMware NSX to address east-west network traffic will help agencies at all levels of government reach their cybersecurity goals.

vmware®

carahsoft.

For more information on VMware NSX, please visit virtualizeyournetwork.com.

DIAGNOSING

Big Da

Government agencies are making strides testing uses of big data to predict risks of disease or the path of a killer virus, but hurdles remain, including linking legacy datasets and setting up common vocabularies.

BY BRIAN ROBINSON

The use of big data to rapidly analyze costs, understand public behaviors and anticipate security threats continues to attract the interest of government agencies that see the technology as a way to gain measurable insights into their most demanding problems.

Nowhere are researchers more active in exploring the uses of big data than in government health care organizations, where data scientists are working toward creating reliable tools for predicting a patient's risk of disease or a virus's path of infection.

To some extent health care programs are an obvious target for big data investment. Agencies already have large databases with years of information on diseases and patient health, and they have an urgent need to provide better and more productive information for researchers, doctors and nurses.

The Veterans Health Administration (VHA), for example, has created several big data analytics tools to help it improve health services to its 6.5 million primary care patients.

The VHA's care assessments needs (CAN) score is a predictive analytic tool that indicates how a given veteran compares with other individuals in terms of likelihood of hospitalization or death. The scores are analyzed by VHA's patient care assessment system (PCAS), which uses these scores and other data to help medical teams coordinate patient care.

The technology has changed the whole approach at the VHA from being purely reactive to one in which patients at the highest risk of being hospitalized can be identified in advance and provided services that can help keep them out of emergency rooms and other critical care facilities, according to Stephan Fihn, director of the VHA's Office of Analytics and Business Intelligence.

While still considered fairly rudimentary tools, the CAN score and PCAS demonstrate that big data predictive analytics can work for large populations.

The agency now needs to "markedly ramp that effort up," Fihn said, and to that end the VHA is working on dozens of predictive models that can be deployed over the next decade. The models will show patients that "this what we know about you, here's what we think you need," he said, and be able to do that in a rapid, medically relevant manner.

BIG DATA, OPEN DATA

Big data tools are also being rapidly developed by the Department of Health and Human Services, a sprawling, 90,000-person enterprise that that both creates and uses data for genomics research, disease surveillance and epidemiology studies.

"There are efforts across the department to try and leverage the data we have," said Bryan Sivak, HHS' chief technology officer.

"At the same time a lot of the datasets we maintain, collect, create or curate can be extended to external entities to help them understand aspects of the HHS ecosystem and try to improve on them, such as with CMS (Centers for Medicare and Medicaid Services) claims data," he said.

One such effort is the OpenFDA project, which essentially took three massive Food and Drug Administration datasets through an intensive cleaning process, Sivak said, and then added an application programming interface (API) so people could access the data in machine-readable ways.

OpenFDA was also linked to other data sources, so that users could access related information from the National Institutes of Health and the National Library of Medicine's MedlinePlus .

The project, which launched as a beta program in June 2014, has already helped to create "a lot of different applications that have the potential to really help reshape that part of the (HHS) ecosystem," Sivak said.

Also within HHS, the National Institutes of Health has committed to several big data programs, including its Big Data to Knowledge (BD2K) initiative. The program, begun in late 2013, is aimed at improving researchers' use of biomedical data to predict who is at increased risk of conditions such as breast cancer and heart disease and to come up with better treatments.

BD2K's goal is to help develop a "vibrant biomedical data science ecosystem," that will include standards for dataset description, tools and methods for finding, accessing and working with

Is the term 'big data' passé?

Is it time to ditch the term "big data"?

While it was useful at one point in encapsulating the idea of exploiting huge volumes of structured and unstructured data, many feel it's past its sell-by date.

On the one hand, said Bryan Sivak, chief technology officer at the HHS, it provides a useful shorthand for referring to things without having to explain it every time. But, at the same time, it "obfuscates the differences in different scenarios."

"To me, when you start to use an overloaded term like big data, it can add some connotation that I don't necessarily intend, or that someone else might interpret because they understand the term differently," he said. "In general, I try to avoid using terms like big data as much as possible."

The term itself is not necessarily the problem, said Tim Hayes, senior director for customer health solutions at Creative Computing Solutions, Inc., because it helps frame the discussion. But there's now a lot of hype around the phrase, which can mask the many challenges of moving, mapping and using data that have to be solved before the expectations generated by the hype can be met.

That hype makes big data a very imprecise term for explaining what is happening in areas such as health and medicine, said Stephan Fihn, director of the VHA's Office of Analytics and Business Intelligence. The previously hyped term in medicine was genomics, he said, and though it's had a dramatic influence in certain areas such as cancer, in his day-to-day practice it's so far had no real effect.

"So, if you talk about big data, at the same time you also have to be very clear about where big data can help (in medicine)," he said. "I prefer to talk about high-level analytics rather than big data, and about streamlining and making more accurate what we do."

- Brian Robinson

datasets stored in other locations and training biomedical scientists in big data techniques.

In October last year it announced grants of nearly $32 million for fiscal 2014 to create 11 centers of excellence for big data computing, a consortium to develop a data discovery index and measures to boost data science training and workforce development. NIH hopes to invest a total of $656 million in these projects through 2020.

While physical infrastructure for computational biomedical research has been growing for many years, the NIH said, as data gets bigger and more widely distributed, “an appropriate virtual infrastructure become vital.”

Vs of big data

According to the National Institute of Standards and Technology, big data consists of extensive datasets that require a scalable architecture for efficient storage, manipulation, and analysis. Commonly known as the ‘V’s’ of big data, the characteristics of data that force new architectures include:

VOLUME – the size of the dataset at rest, referring to both the data object size and number of data objects. Although big data doesn’t refer to any specific quantity, the term is often used when speaking about petabytes and exabytes of data.

VELOCITY – the data in motion, or rate of flow, referring to both the acquisition rate and the update rate from real-time sensors, streaming video or financial systems.

VARIETY – data at rest from multiple repositories, domains or types (from unstructured text or images to highly structured databases).

VARIABILITY – the rate of change of the data from applications that generate a surge in the amount of data arriving in a given amount of time.

VERACITY – the completeness and accuracy of the data sources, the provenance of the data, its integrity and its governance.

FUNDAMENTAL CHALLENGES

There are significant challenges to applying big data to health care, especially with so many legacy datasets to be integrated and shared. Even the use of the term big data can cause confusion.

“Within agencies there are different definitions and types of big data,” said Tim Hayes, senior director for customer health solutions at Creative Computing Solutions, Inc., and a former HHS employee who worked on data analytics there.

“You need to be sure, when mapping data from one database to another, that you can match various labels that are used. Two different agencies might use the term ‘research,’ for example, but they may not be compatible.”

There are “very arcane differences” between what you would assume are fundamental and consistent definitions that turn out not to be consistent at all, Sivak agreed. “It’s a big problem for sure.”

Another barrier is the lack of data scientists capable of working with and understanding the needs of data analytics programs. The solution starts with recognizing that such people are not IT workers, but occupy a niche all their own.

“A lot of what they do is not working with technology, but is in understand-

How a computing powerhouse delivers health care

Health datasets come in many orders of magnitude, but few are as large as the public health big data being gathered and analyzed by computers at the Energy Department’s Oak Ridge National Lab.

About four years ago, the ORNL decided to amass as much public health care data as it could and subject it to the analytics engines of its most powerful computers.

“We were in a unique position with our leadership computing resources and data science expertise, and we saw an opportunity to use health data to discover data-driven insights for better health care quality, integrity and policy,” said Sreenivas Sukumar, a researcher in ORNL’s computational sciences division.

To analyze the datasets, researchers used the Lab’s multicore Titan, the second-most powerful computer in the world, Apollo, an in-memory Urika graph-computer built by Yarcdata, and distributed cloud computing-based machines.

The lab also tapped some of the biggest producers of health related data, including the Cancer Genome Atlas, clinicaltrials.gov, Semantic MEDLINE, openFDA, DocGraph and the National Plan and Provider Enumeration System.

In working with the data, the researchers initially encountered computing silos created by existing information architectures that did not scale to the analytics requirements of the large datasets. Consequently, the lab turned to an approach using graph computing, a scalable computing solution capable of uncovering relationships hidden in the data.

The graph computing almost immediately provided insights into some of the datasets, including feedback on understanding fraud, waste and abuse within the federal health care system, according to ORNL researchers.

In one case, the lab was able to identify a health care provider using multiple identities to bill patients. Another case showed guilt-by-association patterns that highlighted the potential for fraud before the provider began billing.

Georgia Tourassi, director of

ing data," said Brand Niemann, a former senior enterprise architect and data scientist at the Environmental Protection Agency, who now heads up the Federal Big Data Working Group, an interest group of federal and non-federal big data experts.

The fact is, many agencies may already have people with such expertise on staff but don't recognize it. It's a matter of identifying the statisticians that are already working with data and giving them more of a mandate and outlet to mine the agency's data, Neimann said.

Get it right, and the results can be transformative.

ACCURACY COUNTS

Any analysis of big data has limited usefulness if the information in the dataset is not accurate to begin with. Until only recently, VHA's Fihn said he had been skeptical that data analytics could reach the levels of accuracy required for clinical use across the VHA. One reason is that, until just a few years ago, the only data available was from health insurance claims.

"In terms of predictive accuracy we use what we call a C statistic," he said. "A wholly accurate predictive model has a C level of 1.0, and the least accurate has a level of zero. Using (health insurance) claims data, the most accurate level we could get was around 0.65, which is not much better than flipping a coin."

Between 2010 and 2011, however, the VHA brought online a corporate data warehouse that combined clinical data from some 126 different versions of the VISTA (Veterans Health Information Systems and Technology Architecture) electronic health record the agency had been using since the late 1990s.

With that, Fihn said, and greater availability of data on patient medications and vital signs, predictive models are regularly reaching C levels of 0.85, and are pushing 0.9.

It was a "quantum jump" in terms of the usefulness of predictive analytics, he said, and VHA medical staff feel they can now predict with confidence who the high-risk patients are. And even though predictions are still being published using claims data alone, he said, "for our considerations, we now reject those below C levels of 0.85, and we are actually moving to push things as close as we can to 0.9."

HHS doesn't have any global metrics or milestones it wants to reach for big data, Sivak said, though there are specific goals for individual programs. In fact, NIH may have the most expansive set of goals, with BD2K just part of a larger portfolio of activities that NIH is promoting, including cross-agency and international collaboration on big data initiatives and policies.

It's all a marker for just how quickly minds have changed over big data, Sivak believes. "Back in the day," nobody would have given any thought to making datasets public or making them available widely within HHS. But over the past five years the value of that has been conclusively demonstrated, he said, "and as a result, the default setting within HHS has changed from closed to open." •

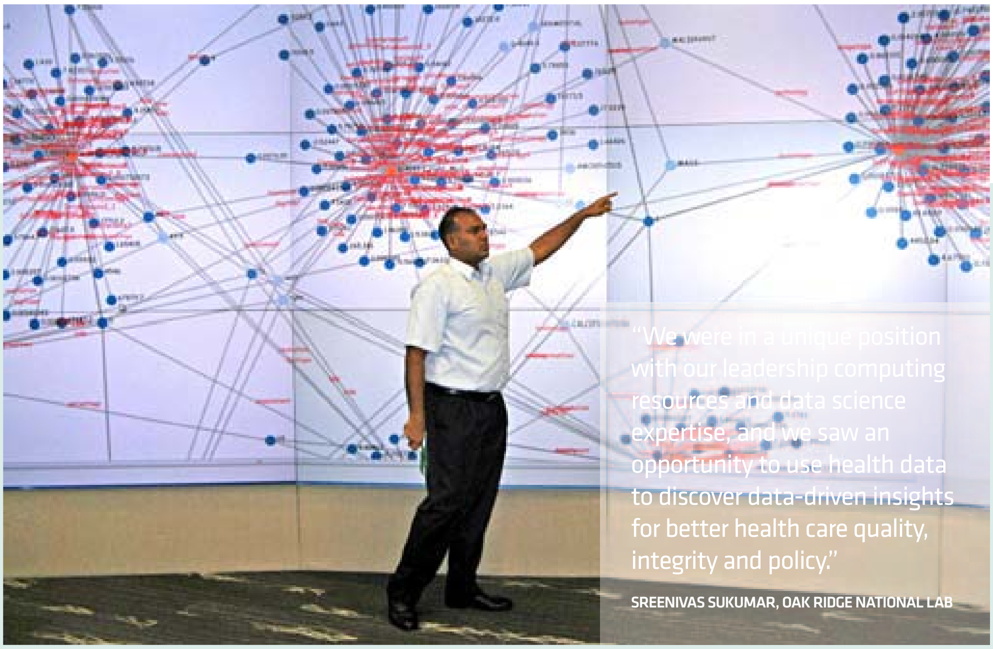

"We were in a unique position with our leadership computing resources and data science expertise, and we saw an opportunity to use health data to discover data-driven insights for better health care quality, integrity and policy."

SREENIVAS SUKUMAR, OAK RIDGE NATIONAL LAB

insights

ORNL's Health Data Sciences Institute (HDSI), said ORNL's approach is novel in health care. Big data computing capabilities in facilities such as ORNL, "are critical to health care delivery," Tourassi said. "It's a paradigm shift in an environment that has always been reactive."

HDSI is reaching out to partners who have different types of data and diverse needs for data analysis – such as genomics, electronic health records and health-sensor data. The projects will help collect, store, integrate and analyze data in support of next-generation personalized medicine care, said the researchers.

For example, ORNL is building the capability for clinical experts to "semantically reason" with medical records and associate health data types, such as claims and clinical records while simulating the outcomes of different clinical interventions.

"We know for certain that health data will be getting bigger and more complex as the practice of medicine expands and progresses, said Tourassi. "By being involved and leveraging the investment, we can anticipate and prepare for the next bottleneck."

The push for geospatial data integration

Agencies are building tools to more easily integrate mapping data with varying formats and legacy sources

BY PATRICK MARSHALL

From NASA to NOAA, and the Department of Agriculture to the Department of Defense, federal agencies and departments are collecting geospatial data at a pace that would choke many server farms. The data isn't just being amassed: It's being analyzed to guide troops in unfamiliar terrain, to track the spread of disease and to decipher crime patterns across the law enforcement enterprise.

While geographic information systems (GIS) have become well established in the federal government, the current challenge for agencies is to develop tools to connect their geospatial programs with their counterparts in other jurisdictions.

While some agencies have been coordinating their data collection and analyses, as a general rule data collected by one agency can rarely be integrated easily and effectively with data collected by another.

Steps are now being taken to change that. In fact, the Federal Geographic Data Committee (FGDC) – the main federal unit charged with integrating federal geospatial efforts – in its National Spatial Infrastructure Plan 2014-2016, said its primary goal is to "develop capabilities for national shared services."

According to agency geospatial leaders and technology experts, there are two major hurdles to achieving those goals: neither the resolution of collected data nor the formats for creating metadata are compatible across data fields.

Not surprisingly, the lead organization researching geospatial data integration and developing tools to enable such integration is USGS. In fact, the agency has created the Center of Excellence for Geospatial Information Science (CEGIS) to lead the effort.

According to Lynn Usery, director of CEGIS, the initial impetus for finding ways to integrate geospatial data was the USGS National Map – a collaborative effort among USGS and other federal, state and local partners to improve and deliver topographic information across the nation. The first hurdle Usery's team faced was the different resolutions of collected data.

"The reason we started the [CEGIS data integration] project was that when we were developing the national map, we realized that the different layers of the national map were all actually compiled and generated separately," said Usery. That meant that when the layers were put together, it might look as if they didn't match up.

Usery's team then tried to determine exactly what it means for a dataset to

CyberGIS: Community-specific infrastructure

CyberGIS is a geospatial-specific infrastructure that manages, processes and visualizes massive and complex geospatial data, while performing associated analysis and simulation.

To drive development of the technology, a consortium of government, academic and private-sector partners has come together to build the National CyberGIS Facility at the University of Illinois, Urbana-Champaign.

With funding from the National Science Foundation, the group will build a high-performance computing system optimized to deal with geospatial data. The system will be equipped with more than 7 petabytes of raw disk storage, solid-state drives, advanced graphics processing units, a high-speed network and dynamically provisioned cloud computing resources.

"There are critical problems that cyberGIS can assist in, from mapping water resources across local, regional and global scales to managing the preparation and response to disasters and emergencies," said Shaowen Wang, the founding director of the CyberGIS Center.

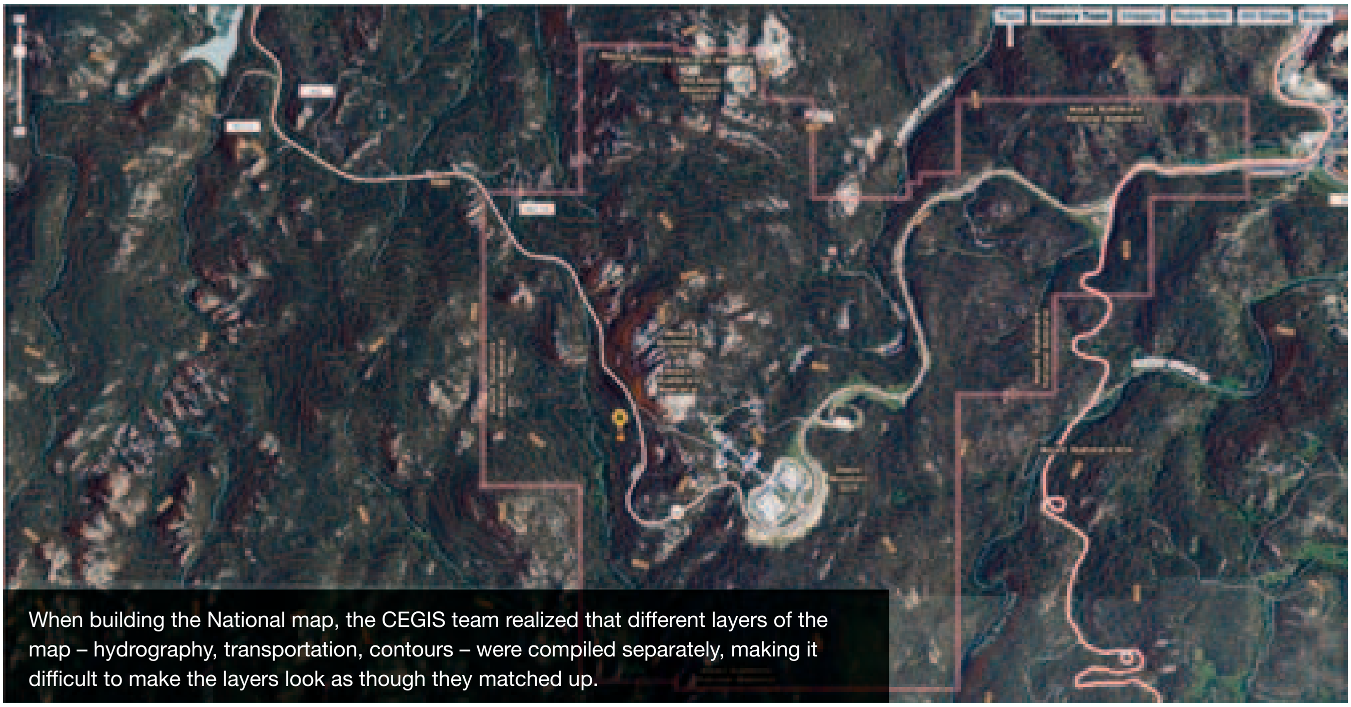

When building the National map, the CEGIS team realized that different layers of the map – hydrography, transportation, contours – were compiled separately, making it difficult to make the layers look as though they matched up.

NATIONALMAP.GOV

be integrated. According to Usery, the team found that if the resolution of two datasets – say transportation data superimposed on image data – was within about 6.4 meters, users would perceive the data as being integrated.

But unless both datasets were already geocoded using the same projection system, getting the two datasets to align correctly can be a major problem. For each data set, an application needs to be created to perform the integration. In the case of integrating transportation data with underlying imagery, said Usery, USGS worked with researchers at the University of Southern California.

In similar fashion, USGS provided support to another group to integrate hydrography data with contour layers in the National Map.

Of course, integrating data sets collected by agencies would be much faster, easier and less expensive, if the data sets were created using standardized metadata – the data about the data, including geographic coordinates and object labels.

To address the problem, CEGIS is considering a “semantic” approach that allows data to be used across application and agency boundaries.

“We are primarily looking at ontology and semantics as a way to integrate data across a variety of organizations and different kinds of data layers,” said Usery. As it is now, government agencies at all levels – as well as the private sector – apply different labels to features, spatial concepts and other data. Accordingly, CEGIS is working to develop a single integrated language or ontology for describing geospatial data.

“We’re taking geospatial data and building an ontology for the data based on all of our features, relationships and interrelationships of features, and then we structure the data using RDF – resource description framework,” said Usery.

“If our data is structured in that form and other data is structured as RDF we can actually bring the data sets together.”

In fact, USGS did just that with data it brought in from the Environmental Protection Agency, which was also in RDF format. “We just ran an automatic query to locate all of the EPA pollution sites within five miles a local firehouse,” said Usery.

DEALING WITH LEGACY DATA

Not surprisingly, the biggest snag to implementing a robust scheme for organizing metadata is the existence of large amounts of legacy data.

According to Usery, USGS has not converted all of its data to RDF. “Our data resides in GIS format and it works very well,” he said. “There are lots of procedures designed around those things and we can’t just completely change over and lose all the legacy developments that we’ve done around GIS platforms.”

Instead, the team has developed a tool that allows analysts to take any section of vector data sets and convert it to RDF.

For now, Usery said, at USGS most efforts to integrate geospatial data continue to be between agencies that share an immediate interest in the specific data. “We try to leverage data from other agencies and not have to collect all the data ourselves,” he said. •

Chicago builds ETL toolkit for open data

The extract-transform-load tool automates the process of extracting information from a database and uploading it to an open data portal

BY STEPHANIE KANOWITZ

Data officials in Chicago are churning out open datasets faster than ever by using technology rather than developers to get the job done.

About a year ago, the city government embedded Pentaho Data Integration (PDI), a graphical extract-transform-load (ETL) tool with pre-built and custom components to process big data, into its Open ETL Utility Kit.

The kit provides several utilities and a framework to help governments extract data from a database and upload it to an open data portal using automated ETL processes.

Before working with PDI, city workers updated datasets manually, said Jon Levy, open data program manager at the Chicago Department of Innovation and Technology. "Most of it was custom-written Java code and that just became difficult to maintain and enhance because it required a very complicated skill set," Levy said.

That also meant Java developers were spending time on updates rather than writing applications that could help city workers and residents, added Tom Schenk, the city's chief data officer.

"What's different now is we have a framework that can be easily used by a lot of people," Schenk said. "I could also give that tool to a number of users around the Chicago, and they'd to be able to program ETLs that are going be easier for them to understand and easier for them to create. It allows us to be more nimble."

For example, the city was able to tap into an application programming interface (API) that monitors water quality at Lake Michigan beaches and uses ETL created with PDI to push the information out hourly.

More recently, an organization used the city's open data to create a map-based app showing where people need permits to park in Chicago, Schenk said.

The two officials also use the information to perform advanced analytics and to merge data to develop predictive models. And every time something is uploaded to the portal, they get email notices. Other features include a log to track what's happening in the portal and a tool that lets users monitor how long it took to run ETLs over time to diagnose problems.

Chicago's toolkit is free to download. To set it up, IT managers must configure about six options to orient it and direct it to their portal. From then on, every time an ETL is needed, a base template exists.

"Probably the most complex thing of all is to write the code to extract data from a server," Schenk said. "At that point, you work with a database administrator, who gives you the SQL code or whatever it might be, you extract it from a server and then you configure about four specific things, such as what dataset is this called."

"The whole workflow is not about custom development, it's about entering simple questions, simple procedures, and that allows you to get online and running," Schenk added.

The result is that datasets can be opened and updated in an afternoon because the toolkit removes the need to wait on staff members. The data portal can also be extended without having to add more people, said Schenk, noting that automation is the key to successful open data programs without large staffs.

> "What we're looking forward to is getting feedback from other cities to see how the ETL can be used and how it can be tweaked."
>
> – TOM SCHENK, CITY OF CHICAGO CHIEF DATA OFFICER

Chicago had all the infrastructure in place when it started. The toolkit is compatible with Windows, Apple and Linux operating systems, making integration easy, Schenk said.

He hopes officials at other city governments use the technology to further their own open data programs.

"What we're looking forward to is getting feedback from other cities to see how the ETL can be used and how it can be tweaked," Schenk said. •

Sowing the Internet of Things into agriculture

Precision agriculture: Information from sensors on farm machinery can increase productivity 10 to 30 percent

BY PATRICK MARSHALL

When it comes to mobile technologies, the Department of Agriculture's National Agricultural Statistics Service (NASS) has already pushed the envelope more than most agencies. Several years ago, NASS rolled out iPads to its 3,500 field enumerators – staff who visit farms around the country to collect information on crop yields, soil conditions and a variety of other data.

The deployment of iPads not only speeded the processing of data, according to Michael Valivullah, chief technology officer at the NASS, it also resulted in savings of $3 million to $5 million in printing and mailing costs.

At the same time, the deployment presented challenges. To secure the data, Valivullah's team configured the iPads so that the collected data was transmitted immediately to the agency's servers rather than stored on the local device. And if there was no Wi-Fi – as is common in the rural areas to which NASS enumerators are sent – the data is cached and automatically sent as soon as connectivity is restored.

While NASS's mobile deployment has already garnered strong returns, Valivullah said he's aiming much higher. "My main interest is to use the Internet of Things in mission-to-mission communications," he said.

Much of the information NASS staff collects – how much land was sown with what seeds, how much water was used, how much was harvested, etc. – is automatically being collected by sensors carried by sophisticated farm machinery such as harvesters. "With the Internet of Things, we would be able to get the data from these onboard systems," said Valivullah.

Merge that data with satellite data already being collected by the government, and the department, Valivullah said, will be able to make unprecedented contributions to precision agriculture, which lets farmers (or their machinery) determine the exact location where crops need to be irrigated or fertilized to maximize harvest. "By applying precision agriculture and big data we can increase productivity by another 10 percent to 30 percent and make farmers more profitable," he said.

> "With sensors on combines, we could get resolution down to 1 centimeter."
>
> – MICHAEL VALIVULLAH, CTO AT UDSA'S NATIONAL AGRICULTURAL STATISTICS SERVICE

BETTER RESOLUTION CREATES MORE DATA

The key value of using combine sensors to collect data is resolution. "Satellite information has a resolution of about 100 square feet," said Valivullah. "With sensors on combines, we could reduce that down to a resolution of 1 centimeter."

There are two primary challenges Valivullah sees in trying to integrated combine sensor data. The first is getting access to the data. In fact, who owns the data is not clear, especially when it is collected by sensors on rented machinery. At least one private company – Climate Corp., which was purchased by Monsanto in 2013 – is looking to use combine sensor data to deliver consulting services to farmers.

But even if the government secures access to the data, Valivullah sees challenges in managing the huge amounts of data that are collected. "Right now, to produce one cropscape map we're taking more than 100 billion pixels," said Valivullah. "Multiply that another 10 times and we're talking about trillions of pixels. Do we really need to collect all that information? And if we really need it, how are we going to filter it?"

According to Valivullah, his agency has teams that meet every two weeks to generate plans for integrating the data from satellites, mobile sensors and human enumerators. "Hopefully, we will be able to design some sort of data handling methodology to deal with the big data that is going to be coming down the pike," he said. •

Visual Studio LIVE!

EXPERT SOLUTIONS FOR .NET DEVELOPERS

vslive.com/austin

Austin JUNE 1-4

HYATT REGENCY, AUSTIN, TX

DON'T MESS WITH CODE

Visual Studio Live! is pullin' into Austin in June, where the Texas attitude is big, and the code is bigger! For the first time in 8 years, Visual Studio Live! is bringing its unique brand of practical, unbiased, Developer training to the deep heart of Texas. From June 1 - 4, we're offering four days of sessions, workshops and networking events—all designed to make you better at your job.

DEVELOPMENT TRACKS INCLUDE:

- Visual Studio / .NET
- ASP.NET
- JavaScript / HTML5 Client
- Cloud Computing
- Mobile Client
- Database and Analytics
- Windows Client

Register NOW and Save $300!

Scan the QR code to register or for more event details.

Use promo code VAUMAR2

vslive.com/austin

GOLD SPONSOR

GitHub

SUPPORTED BY

Microsoft Visual Studio msdn magazine Visual Studio MAGAZINE

PRODUCED BY

1105 MEDIA

ADVERTISER INDEX

This index is provided as an additional service. The publisher does not assume any liability for errors or omissions.

MEDIA CONSULTANTS

Mary Martin
(703) 222-2977
mmartin@1105media.com

Bill Cooper
(650) 961-1760
bcooper@1105media.com

Matt Lally
(973) 600-2749
mlally@1105media.com

Ted Chase
(703) 876-5019
tchase@1105media.com

PRODUCTION COORDINATOR

Lee Alexander
(818) 814-5275
lalexander@1105media.com

GCN HAS GONE MOBILE.
Go to gcn.com/tablet and download the tablet app today!
Your mobile gcn.com experience — optimized.
Visit gcn.com from your smartphone and enjoy the easier navigation and new sharing options

PUBLIC SECTOR
MEDIA GROUP

CORPORATE HEADQUARTERS
9201 Oakdale Ave., Suite 101
Chatsworth, CA 91311
www.1105media.com

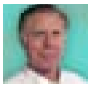

EMERGING TECH

BY PATRICK MARSHALL

A network for the Internet of (moving) Things

WHAT COULD MAKE more sense than to make mobile networks, and not just client devices, mobile?

That's the idea behind Veniam, a company that not only collects data from sensors – whether stationary or deployed on fleets of vehicles – but also delivers Wi-Fi connectivity to those within range of the vehicles. Initial deployments are being targeted at fleets in metropolitan areas – such as transit or garbage collection vehicles – and secure areas, such as ports.

"As a professor I've been working on connected vehicles since 2005," explained João Barros, co-founder of Portugal-based Veniam. "I was mostly intrigued by how one could use vehicles as mobile sensors to get as much data as possible about a city. That's the way I got involved in connected vehicles research – as a way to build an urban scanner."

At the same time, Barros, the former national director of the Carnegie Mellon Portugal program, and his partner, Susana Sargento, a professor at Portugal's University of Aveiro, realized such a system would be an easier sell if it was delivering service the other way, too, by providing Wi-Fi connections to users on or within range of the vehicles.

The team also realized that, whether it's sensors or cell phones, they needed to find a middle ground between the expensive but far-reaching cellular communications and the low-cost but short-range of Wi-Fi.

The solution Veniam came up with was part hardware, part software and part process.

The hardware piece of the puzzle is called NetRider, a box that's about the size of a thick book. The boxes are not only Wi-Fi hotspots, they also include either 3G or 4G cellular interfaces, all built out of off-the-shelf components.

Cities can deploy sensors that don't have Internet connections and when a NetRider-equipped vehicle goes by, the data will be picked up.

But the special sauce in NetRider is in the software, which scans available cellular and Wi-Fi connections and routes data in the most efficient way, ensuring that both end-user data – even streaming YouTube – and sensor data is not dropped as the vehicle moves from one network connection to another.

While NetRider supports Wi-Fi connections from devices employing all the standard IEEE protocols, Veniam engineers have also developed their own routing protocols that have enabled greater transmitting range. According to Barros, the NetRider can provide signal to devices as much as 1,000 meters away where there is a line-of-sight connection and 300 to 400 meters in typical urban environments.

The system has already been deployed in Porto, Portugal, on a fleet of more than 100 buses and other transit vehicles.

"Initially we use the cellular backhaul to learn precisely where most of the traffic is being consumed, and then we use this information to deploy access points for infrastructure communications, where you need most of the traffic to be offloaded from cellular," said Barros.

While the benefits of the mobile Wi-Fi are obvious for transit riders, the benefits of creating networks moving data from sensors are at least as great. "We also collect data from sensors, sensors that are on the vehicle and sensors that are outside of the vehicle but that use the vehicles as data couriers," said Barros.

Cities can deploy low-cost sensors that don't have Internet connections – sensors monitoring air quality, noise, or even how full a garbage container is – and when a NetRider-equipped vehicle goes by, the data will be picked up. "When a public bus, taxi or garbage collection truck comes by, it syncs with the sensor, it gets the data, stores the data and whenever the vehicle is within range of an access point it sends the data to the cloud," said Barros.

Veniam's mobile networks offer both end-to-end encryption and authentication. Authentication is generally not required on public buses. "Our customers opted not to have logins," said Barros. "So it works without authentication, and we are not tracking any user behavior."

According to Barros, the NetRider technology is especially appropriate for ports, container terminals and even military bases, where cellular doesn't work well because large metal containers make the signal propagation less efficient. "Normal Wi-Fi has a range that is too small for such spaces," said Barros.

"Our technology allows vehicles to communicate over ranges that are 10 times larger than traditional Wi-Fi. We also were able to establish connections in 2 milliseconds, as opposed to a few seconds that normal Wi-Fi requires." •

GCN

Technology,
Tools and Tactics
for Public Sector IT

Where you need us most.

GCN
Technology, Tools and Tactics for Public Sector IT
GCN Newsletter
Sign up for our newsletter.
Email Address
Agencies must update to newer versions of Transport Layer Security
DENIAL-OF-SERVICE
DEFENSE

Mobile

Tablet

Desktop

Print

SANS

The Most Trusted Name in Information and Software Security Training

SANS 2015
will be held at the

Walt Disney World Swan and Dolphin Resort

"You're getting training from instructors who do this stuff for a living and not from someone who just teaches."

-Scott Avvento, Alpine Cyber Solutions, LCC

To learn more about the SANS 2015 cybersecurity training event, or to register, please visit:

www.sans.org/u/QM

2015

Orlando, FL

April 11-18, 2015

SANS' Most Comprehensive Training Event of the Year…something for everyone!

More than 35 World-Class Information Security Courses

Top-Rated SANS Instructors

New Courses Including
SEC511: Continuous Monitoring and Security Operations

CORE NetWars & DFIR NetWars Tournaments

GIAC Approved Training

Save $200
Register & Pay by March 18th