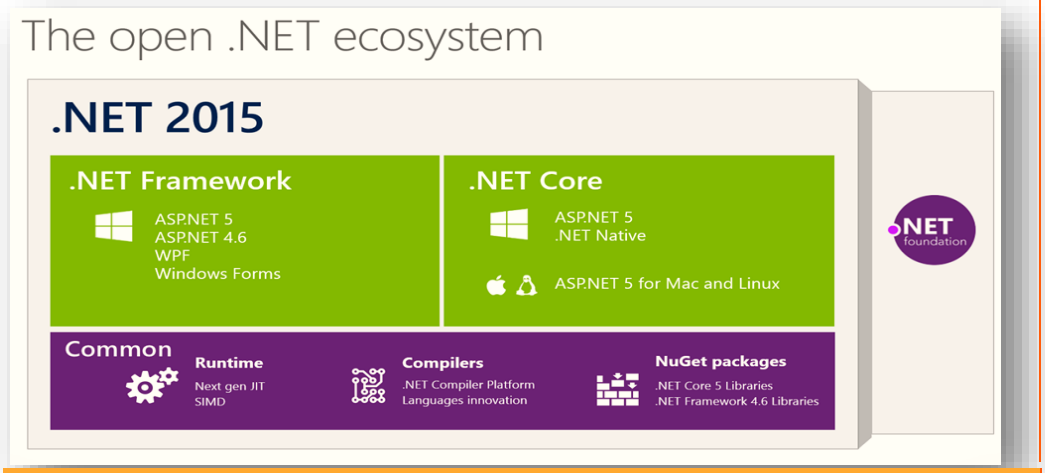


# TechTalk

## Issue 30

## December 2014

LISTENING TO BUSINESS, APPLYING TECHNOLOGY

IN THIS ISSUE



# Microsoft Open Sources .NET Core, Runtime And Its Framework Libraries

Microsoft announced that .NET Core will be open source, including the runtime as well as the framework libraries. .NET Core is a modular development stack that is the foundation of all future .NET platforms. It's already used by ASPNET Core 5 and NET Native. Microsoft decided to open source .NET Core to lay the foundation for a cross platform .NET and build and leverage a stronger ecosystem.

"Today, people who are stuck on the .NET platform have to use a server environment that doesn't have what Linux does," says James Watters, who, at a company called Pivotal, works hand-and-hand with a wide range of developers and companies as they build large online software applications. "They're stuck with a generation-behind technology." For Watters, Microsoft has ample ground to make up. But in opening sourcing what's called the .NET Core runtime—freely sharing it with the world at large—the company at least gives itself a fighting chance as it seeks to maintain a hold on the way the world builds and runs software.

In theory, an open source .NET that runs on Linux and Mac OS will expand the use of Microsoft's developer tools. Then the company can pull in revenue through other channels—through premium versions of its developer tools and through its cloud computing service, Microsoft Azure.

Among developers and businesses building websites and other large online services, .NET is one of the primary competitors to Java. It's widely used among companies that rely heavily on Microsoft software —the company says .NET was installed more than 1.8 billion times over the last year.

"This is great," Watters says of an open source .NET, "but it would be even more powerful if they made their operating system open source too. Those .NET runtime facilities interface with low-level libraries in the OS. If the entire stack is open source, folks can really optimize things." Yes, he's calling on Microsoft to open source its crown jewels: Windows.

## Panduit Introduces No Battery, Maintenance-Free Uninterruptible Power Supply (UPS)

Panduit Corp., a global leader in Unified Physical Infrastructure SM (UPI)-based solutions, announces the launch of the Industrial Network Uninterruptible Power Supply (UPS). The Industrial Network UPS lowers the risk of downtime due to battery failures by 39% to keep key equipment in facilities operating at full capacity. It is designed specifically to back up industrial managed Ethernet switches.

With many manufacturers and processors dealing with the issue of a shortage of skilled labor, maintaining batteries in UPS systems can often fall through the cracks, and often goes unnoticed until the UPS is needed, and by then it is too late. According to Paul Schildhouse, Senior Product Manager, Panduit Corp., the Panduit Industrial Network UPS uses ultra-capacitor technology instead of batteries, eliminating the number one cause of conventional UPS failures, which lowers the risk of downtime. The no battery design is maintenance-free, providing 2X greater ROI and 50-70% lower cost of ownership than an ordinary UPS.

# The Future is Now

## High-tech farming poised to change the way the world eats

Investors and entrepreneurs behind some of the world's newest industries have started to put their money and tech talents into farming -- the world's oldest industry -- with an audacious and ambitious agenda: to make sure there is enough food for the 10 billion people expected to inhabit the planet by 2100, do it without destroying the planet and make a pretty penny along the way.

Silicon Valley is pushing its way into every stage of the food-growing process, from tech tycoons buying up farmland to startups selling robots that work the fields to hackathons dedicated to building the next farming app.

Dozens of companies are creating technology to make farmland more productive and farming more efficient, using robots to trim lettuce or software to calculate grass production for cattle grazing. Others are tapping technology to find substitutes for meat, cheese and eggs, so less land is used to raise livestock, fewer greenhouse gas-spewing trucks are used to transport them, and fewer animals are subject to inhumane slaughter. VCs have propped up startups such as Hampton Creek, which sells mayonnaise and cookies that use plant products instead of eggs, and Impossible Foods, a Redwood City company making hamburgers and cheese without meat or dairy.

And Silicon Valley isn't just making technology for farms -- some of its highest-profile investors are buying farmland to have a hand in how farmers work their fields and influence the type of food that's available for future generations of consumers.

Farmers, for the most part, have welcomed Silicon Valley techies into their world of planting seasons and water woes. From Salinas to Fresno to Tracy, California farmers are increasingly tech-minded, running their fields from iPads and tracking soil moisture and nitrogen levels with cloud software programs, and they are hungry for more tech solutions that will bolster their land's productivity and their bank account balance.

## Ship of the Future

*Crave's Michael Franco sets sail on Royal Caribbean's latest ship and gets treated to a tour of its high-tech features from virtual balconies to power robots.*

***The Quantum of the Seas, a tech-rich beauty of a ship***

The numbers are certainly impressive. It weighs 168,666 gross tons. It's 2.5 times taller than the great Pyramid of Giza. It's longer than five Boeing 747 jetliners and 11 times bigger than a blue whale. And its four thrusters produce 4,694 horsepower each -- the same as 24 Formula One race cars.

But numbers can only get you so far. I was eager to see just how impressive Royal Caribbean's new **Quantum of the Seas ship** was in person, so when members of the media got invited aboard for a recent two-night sail out of New York, I jumped on deck.

Billed as the world's first "smart ship," the Quantum tucks cutting-edge technology into every nook from bow to stern. There's a bar where robotic arms mix your drinks; a theater where six "robo screens" star in a tech-heavy performance; 30 bumper cars; an Xbox playroom; a hydraulic arm with a glass-enclosed capsule dangling from it that lifts you 303 feet above the sea; and a skydiving simulator that lets you fly on a column of air in the middle of the ocean.

Fun aside, perhaps the most useful aspect of the ship's technology is its fast Wi-Fi. As I reported back in August, Royal Caribbean forged a partnership with O3B to launch its own set of low-Earth-orbit satellites to provide faster connectivity to its ships. By beaming signals to dedicated satellites three to four times closer to Earth than those used by previous systems, Royal Caribbean says it now has more bandwidth than all the other ships at sea combined.

Indeed, lots of people passed the entire cruise with their faces buried in their smartphones and tablets (maybe an unpleasant side effect of providing all that bandwidth to people who are supposed to be relaxing), so the system was definitely put to the test. Royal Caribbean says the onboard speeds should be similar to those on land, and I found this mostly true, with some brief slowdowns here and there. Interestingly, the O3b network doesn't turn on till you're out at sea, so things got better once we were away from New York harbor.

Quantum of the Seas will sail out of Bayonne, N.J., and make trips through the Caribbean until May 2015. After that, it will be relocated to China, so if you want to take a trip aboard the world's first smart ship without paying a boatload l \n airfare, now's the time. Cruises are available starting at $499 per person for a "Three-Night Sampler Cruise" that goes along the New Jersey coast, to a 7-night Bahama cruise near the holidays starting at $2349 per person.

*(Reported by Michael Franco-CNET)*

# Technology Focus

## Deduplication best practices and choosing the best dedupe technology

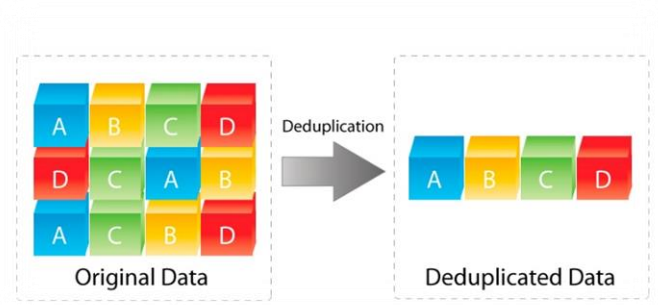


Data deduplication is a technique to reduce storage needs by eliminating redundant data in your backup environment. Only one copy of the data is retained on storage media, and redundant data is replaced with a pointer to the unique data copy. Dedupe technology typically divides data sets in to smaller chunks and uses algorithms to assign each data chunk a hash identifier, which it compares to previously stored identifiers to determine if the data chunk has already been stored. Some vendors use delta differencing technology, which compares current backups to previous data at the byte level to remove redundant data.

Dedupe technology offers storage and backup administrators a number of benefits, including lower storage space requirements, more efficient disk space use, and less data sent across a WAN for remote backups, replication, and disaster recovery.

While the overall data deduplication concept is relatively easy to understand, there are a number of different techniques used to accomplish the task ofeliminating redundant backup data, and it's possible that certain techniques may be better suited for your environment.

*Source Deduplication vs. Target-* Deduping can be performed by software running on a server (the source) or in an appliance where backup data is stored (the target). If the data is deduped at the source, redundancies are removed before transmission to the backup target. If you're deduping right at the source, you get the benefit of a smaller image, a smaller set of data going across the wire to the target. Source deduplication uses client software to compare new data blocks on the primary storage device with previously backed up data blocks. Previously stored data blocks are not transmitted. Source-based deduplication uses less bandwidth for data transmission, but it increases server workload and could increase the amount of time it takes to complete backups.

**Source deduplication is well suited for backing up smaller and remote sites** because increased CPU usage doesn't have as big of an impact on the backup process. Virtualized environments are also "excellent use cases" for source deduplication because of the immense amounts of redundant data in virtual machine disk (VMDK) files. However, if you have multiple virtual machines (VMs) sharing one physical host, running multiple hash calculations at the same time may overburden the host's I/O resources.

Most well-known data backup applications now include source dedupe, including Symantec Corp.'s Backup Exec and NetBackup, EMC Corp.'s Avamar, CA Inc.'s ArcServe Backup, and IBM Corp.'s Tivoli Storage Manager (TSM) with ProtecTier.

Target deduplication removes redundant data in the backup appliance -- typically a NAS device or virtual tape library (VTL). Target dedupe reduces the storage capacity required for backup data, but does not reduce the amount of data sent across a LAN or WAN during backup.

**A target deduplication solution is a purpose built appliance, so the hardware and software stack are tuned to deliver optimal performance.** So when you have large backup sets or a small backup window, you don't want to degrade the performance of your backup operation. For certain workloads, a target-based solution might be better suited.Target deduplication may also fit your environment better if you use multiple backup applications and some do not have built-in dedupe capabilities. Target-based deduplication systems include Quantum Corp.'s DXi series, IBM's TSM, NEC Corp.'s Hydrastor series, FalconStor Software Inc.'s File-interface Deduplication System (FDS), and EMC's Data Domain series.

*Inline Deduplication vs. Post-processing Dedupe-* Another option to consider is when the data is deduplicated. **Inline deduplication removes redundancies in real time as the data is written to the storage target.** Software-only products tend to use inline processing because the backup data doesn't land on a disk before it's deduped. Like source deduplication, inline increases CPU overhead in the production environment but limits the total amount of data ultimately transferred to backup storage. Asigra Inc.'s Cloud Backup and CommVault Systems Inc.'s Simpana are software products that use inline deduplication.

**Post-process deduplication writes the backup data into a disk cache before it starts the dedupe process.** It doesn't necessarily write the full backup to disk before starting the process; once the data starts to hit the disk the dedupe process begins. The deduping process is separate from the backup process so you can dedupe the data outside the backup window without degrading your backup performance. Post-process deduplication also allows you quicker access to your last backup.

*Global Deduplication-* **Global deduplication removes backup data redundancies across multiple devices** if you are using target-based appliances and multiple clients with source-based products. It allows you to add nodes that talk to each other across multiple locations to scale performance and capacity. Without global deduplication capabilities, each device dedupes just the data it receives. Some global systems can be configured in two-node clusters, such as FalconStor Software's FDS High Availability Cluster. Other systems use grid architectures to scale to dozens of nodes, such as Exarid Systems'DeltaZone and NEC's Hydrastor.

The more backup data you have, the more global deduplication can increase your dedupe ratios and reduce your storage capacity needs. Global deduplication also introduces load balancing and high availability to your backup strategy, and allows you to efficiently manage your entire backup data storage environment. Users with large amounts of backup data or multiple locations will gain the most benefits from the technology. Most of the backup software providers offer products with global dedupe, including Symantec NetBackup and EMC Avamar, and data deduplication appliances, such as IBM's ProtecTier and Sepaton's DeltaStor offer global deduplication.

# Tech News

## New IBM service shines a light on mobile device and app performance

With mobile devices rapidly becoming the tools of choice for enterprise work, IBM wants to help IT departments make sure they can serve all users. One new service helps IT departments gauge how applications on mobile devices are working, while the other offers a way to deliver them virtually. Both are available now and work

on any major mobile OS, said Linda Lyding, director of portfolio strategy and development.

IBM Mobile Infrastructure Analytics Services can show how well mobile applications are performing, both in terms of speed and the user experience. It's designed to help CIOs make decisions about enterprise IT infrastructure and work with software development teams to make sure users are getting what they need out of each application. It's packaged as SaaS (software as a service).

The service can gather data in real time from the company's network and servers and the mobile devices themselves. A dashboard presents that data so IT can act on it. This provides background information such as what kinds of devices a company's employees are using. It also can show information whether an app is running slowly on users' devices and why that might be, including shortfalls in back-end computing, bottlenecks in the LAN or problems with how the code runs on devices.

The user-experience end of the service is powered by IBM's Tealeaf CX Mobile, a tool that's also used to manage the customer experience on apps offered to the public. Here, it's dedicated to internal app use by employees. Tealeaf can tell development teams immediately about problems users may have in using an app.

Companies can monitor all this after an app has been released to all users or during the testing phase.

## Cisco Unveils Enterprise Collaboration Tool

Building on its previous work in electronic collaboration forms that include audio and video conferencing, Cisco Systems is now offering ways to manage workers' day-to-day interactions and tasks. The networking giant announced the availability of Project Squared, a service aimed at offering an online gathering place for getting work done.

Project Squared is delivered via an app or the Web. It is the latest response to changes in the way people work in 2014, which often involves being on the move with smartphones and tablets instead of in front of computers, according to Cisco. These mobile workers are more at home using instant messaging and find e-mail cumbersome. That development has encouraged a wave of next-generation collaboration tools.

Project Squared combines communications that take place in real time, such as videoconferences and meetings carried via Cisco's WebEx services, with those that don't such as posted messages, documents and presentations.

One of the main selling points of Project Squared is that it gives teams a place to work together virtually, according to Cisco. In one scenario, teams can download Project Squared, start a room on any topic, and invite others inside or outside the organization to join by entering their names or e-mail addresses. Once in the room, team members can securely share messages, post files, and start voice and video calls. Files created in the cloud are viewable immediately alongside messages, with no download necessary. It solves the problem with real time connecting to non-real time.

## Yahoo and Flickr mix social computing, geolocation and computer vision to boost image recognition tech


Not long ago, the concept of image recognition was considered the wave of the future, somewhere on the sidelines of public consciousness. Not anymore.

Photography trends have now propelled this technology to front and center stage because of the sheer volume of images shot and stored and the need to find them quickly and easily. And as technology advances, image recognition joins other techniques to enhance search algorithms and results. For example, Yahoo's Flickr and Labs teams have already developed a viable search algorithm using a multi-pronged strategy of computer vision, geographic information and human interaction to move beyond the traditional reliance on metadata. Yahoo's Project Weather, which powers the Yahoo Weather app, offered an opportunity for the company to find and serve location-based images on request. It gathered metadata, geolocation and social interactions of Flickr photos, combined with images hand-picked by editors, to select the most interesting photos for the app.

Instead of relying exclusively on user-generated content such as tags and photo titles, Yahoo turned to computer vision and deep learning to identify objects in images. This combination of social and computer vision increased the quality of search results. Editors searched for photos from the service, Flickr members submitted their photos or were contacted for permission to use their photos, and Yahoo Labs analyzed the images based on social media interactions such as comments, likes and favorites to showcase the most appealing ones based on human judgments.

That social computing method yielded about 6 million weather photos. After corrections were made for inaccurate geolocations and time stamps, low resolution and errant content, about 1.5 million photos remained in the pool.

# Tech News

## Lenovo cools the Thinkpad Helix with a tiny vapor chamber, no fans necessary

The second-generation Lenovo ThinkPad Helix is one of the first convertible ultrabooks to run Intel's Core M processor, which allows it to do a pretty neat trick. Instead of building a few whirring fans into the case like you'd find on most devices, Lenovo designed a small vapor chamber in the device that cools it without moving any parts. It's also thinner and lighter than the original Helix.

The Helix 2 is essentially a tablet that docks to a keyboard when you need to get some work done. All the important hardware is in the screen portion, which frames an 11.6-inch 1080p touchscreen with enough battery capacity to run for a few hours on its own. That goes to a full 12 hours when docked to the keyboard. It is 0.38-inches thick, which is about the same as the old iPad 4 before the first Air came out. It's more toward the computer end of the weight scale, though, at 1.8lbs.

The vapor chamber that handles all the cooling is less than half a millimeter thick, and completely does away with the need for fans or opening in the casing. That makes the overall device more sturdy and usable. You don't have to worry about blocking the intakes or getting dust built up in a fan. A vapor chamber like this one is becoming more common in enthusiast PC cooling solutions because of its increased efficiency. As the CPU heats up, coolant molecules in the vapor chamber change phase and diffuse outward. When the coolant condenses on cool surfaces, the heat is dissipated away from the CPU. The Helix 2 starts at $999, which is a much more palatable price point than the first-generation model.

# Special Focus

## F5 intros Silverline, a layer 4-7 networking as a Service Platform

F5 Networks Inc. introduced Silverline, a new platform for Layer 4-7 networking as a service. Through Silverline, enterprises will be able to subscribe to Big-IP network services as SaaS-based applications. Silverline is based on F5's acquisition of Defense.Net, a cloud-based distributed denial of service mitigation provider. For the Silverline launch, F5 has expanded the underlying infrastructure of Defense.Net and extended the DDoS protection service portfolio that company offered. It also plans to add other application delivery services over Silverline.

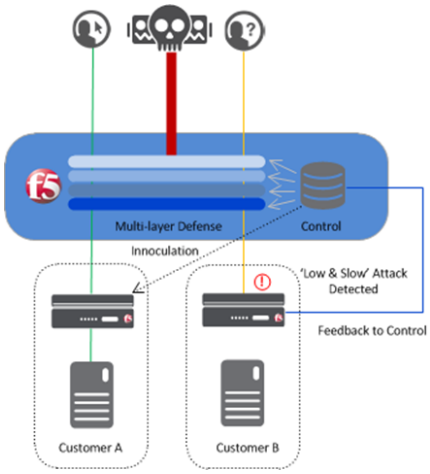


"We're going to expand [Silverline] to things like Web application firewall [WAF] as a service. We're looking at other things like global traffic management and DNS-type services," said Karl Triebes, F5's executive vice president of product development and chief technical officer. "Ultimately we'll be able to offer a large portfolio of our application services that we provide [today] on-premises as a service [later] as part of Silverline."

Security services, like a secure Web gateway, are the best initial use cases for a networking-as-a-service platform like Silverline, said Andrew Lerner, research director for Gartner Inc. However, as enterprises move more applications to the cloud, they will want to consume additional Layer 4-7 services in this manner, he said.

### *Silverline offers alternative way to consume F5 services in the cloud*

The vendor offers Big-IP Virtual Edition via several infrastructure-as-a-service (IaaS) providers, including Amazon Web Services and Microsoft Azure.

One advantage with a platform like Silverline is the reduced operational complexity. Enterprises that consume Big-IP via an IaaS provider need a lot of the same expertise they would need to deploy an F5 appliance in a data center. "You're still managing and running it like an application delivery controller that you own and operate," Lerner said. "If you're running it in the [Silverline] cloud, it's supplementing your on-premises software. In essence [Silverline] is additional services layered in."

"As we go forward, we're able to link these services with the on-premises capabilities that are in our core F5 products," Triebes said. Essentially, F5 customers can use the company's Synthesis architecture, which orchestrates multiple F5 platforms into a fabric, to integrate Silverline services into its other F5-based technologies, whether those are Layer 4-7 services within their own data centers or virtual appliances deployed in the public cloud.

### *F5 Silverline DDoS offers three service models*

The original Defense.Net service will now be known as Ready Defense. F5 added the "Always On" service, which will serve as a data center operator's first line of defense against DDoS attacks, he said. This service will continuously scrub all incoming traffic before it reaches the customer's infrastructure. F5 also introduced "Always Available," a stand-by service that will activate as soon as the customer identifies an attack.

# About Galaxy

- One of the most respected Information Technology integrator of the best of breed products and solutions for Enterprise Computing, Storage, Networking, Security, Automation, Application Delivery, ERP and Business Intelligence.
- An ISO 9001:2008 organization , founded in 1987
- Committed team of over 200 skilled professionals
- PAN India presence
- Trusted IT services provider to more than a 1000 companies
- Experienced consultants certified on a wide spectrum of technologies
- The Galaxy Technology Innovation Centre, a state-of-the-art integrated hardware and software laboratory, allows customers a hands-on look at the latest storage, backup, security, application delivery and virtualization technologies.
- Customer list includes many of India's leading corporations, banks and government agencies
- Four business units collaborate to provide a full spectrum of services and ensure smooth projects. Together, they provide our customers with truly end to end professional IT Services.

**Galaxy Business Solutions**
System integrators of best of breed technologies to deliver solutions to the problems and challenges that confront enterprises

**Galaxy Technology Services**
Skilled pool of resources consistently maintains and delivers enterprise class service levels

**Galaxy Network Solutions**
One of India's most trusted active and passive networking specialists

**Galaxy BI Consulting Services**
Helps organizations to deliver and leverage business intelligence to create substantial business impact

**NEWSLETTER COMPILED BY**

**Galaxy Office Automation Pvt. Ltd.**
A-23/24, Ambika Towers , Ground Floor, Off Jijamata Road, Nr. Pump House, Andheri (E),Mumbai – 400093, India

**Phone: 91-22-42187777**
**Fax: 91-22-421877760**
**E-mail: galaxyinfo@goapl.com**
www.goapl.com

**VISION**

***"To become the most preferred technology solution partner by listening to our customers, anticipating their needs and providing reliability, flexibility, responsiveness and innovative products and services. Achieving market leadership and operating excellence in every segment of our company."***

**MISSION**

***"Total customer satisfaction; through innovative insights, quality service and excellence in technology deployment."***

**VALUE PROPOSITION**

***"We understand the need of a common vendor for all your IT needs. Hence, we are committed to long-term partnerships by delivering on our commitments."***

## MD Speaks

*"As we ring out 2014 and bring in 2015, I want to share my thoughts about the 3 trends that would define the technology industry next year.*

*Mobile computing will continue to lead the charge as it has done for most of this year. All industries, whether old or new, will have to continue investing in this space. Mobile devices will continue to penetrate at rates much higher than other end devices and new applications for traditional business or even new businesses that rely solely on mobile connectivity will look to capitalize on this.*

*Analytics will play an extremely important role as the volume of data generated by the 'internet of things', social media, digitization and availability of public data just keeps on increasing. Enterprises will look to leveraging these sources along with their enterprise generated data to be able to spot trends or predict the future to stay ahead of competition. The buzz that big data created in the past couple of years will become commonplace next year. Companies that specialize in generating and selling data for use by others will spawn and thrive.*

*Security, as always, will be critical the next year too. However, there will be a shift from perimeter based security to application based self-protection. Security-aware application design, static and dynamic application security testing, and runtime application self-protection have all become necessities and this will lead to new models of building security directly into applications. This will require considerable investments in revamping existing applications.*

*Of course, there will be exciting things happening in smart machines and 3D printing but I think those will really be poised for take-off in either 2016 or 2017. Of course, I will be the happiest person if I'm wrong on that count !"*