Unmatched Global Market Coverage

Telecom COTS World
Broadband Broadcast IoT AI Convergence

Telecom IT - HPC - A.I. - IoT
Infrastructures - Networks - Edge - RAN
Cloud - Data Centers - Storage
Video Networks - Broadcast - Digital TV
From Chips to Rack Scale Systems

Headlines
• IASM - Intelligent Application & Service Monitoring Software from 13 Top Vendors
• Radisys: Communication & Digital Engagement Solution for TELCOs & Enterprises
• Wind River: Kubernetes-Based Cloud Native Solution for Complex 5G vRAN Edge
• Kudelski Security & Fortinet: Security Device Support & Management Services
• AKAMAI recognized as a Leader for CDN in 2019 among the Top 8
• Enea Unveils Complete 5G Data Management Suite
• Data Integration Tools: Top 16 Vendors

Tech Giants' Billion-Dollar Acquisitions

- Microsoft
- Facebook
- Amazon
- Google
- Apple

LinkedIn $26.2B

WhatsApp $22.0B

WHOLE FOODS MARKET $13.7B

MOTOROLA $12.5B

skype $8.5B

NOKIA $7.2B

GitHub $7.5B

AQUISITION DATE

2010   2015   2020

Products & Technology from Silicon to Application-ready Platforms

Cover Story:

Visualizing Tech Giants' Billion-Dollar Acquisitions
Over the last 3 decades, the FAMGA tech giants —
Facebook, Amazon, Microsoft, Google, and Apple — have
collectively made over 750 acquisitions.

Headlines:

**• IASM - Intelligent Application & Service
Monitoring "The Top 13 Vendors"**
Today's Applications And Services Require Intelligent
Monitoring

**• Radisys Unveils Engage@Work: A Communication
and Digital Engagement Solution for Service
Providers and Businesses**
One unique application, with multiple use cases – from
customer and brand engagement to employee
collaboration – securely and cost-effectively

**• Wind River Introduces Kubernetes-Based Cloud
NativeSolution for Complex 5G vRAN Network Edge
Needs**
Wind River Cloud Platform combines a fully cloud native,
Kubernetes and container-based architecture with the
ability to manage a physically and geographically
separated infrastructure for vRAN and core data center
sites

• Kontron ATX Motherboards Designed by Fujitsu
BIOS Updates support of the latest Intel® processors for
better performance and longer lifetime

**• HPE advances the cloud experience through
intelligence and composability**
HPE announced expanded intelligence and composability
offerings by integrating HPE Primera storage HPE
Synergy and HPE Composable Rack

**• Software-Defined and Cloud-Native Foundations
for 5G Networks**
Communications service providers (CSPs) are turning to
the adoption of cloud data center innovation both
adapted and re-developed for carrier networks

**• Enea Unveils Complete 5G Data Management
Suite**
Cloud-native 5G data suite to support automated
network processes and data monetization

**• Kudelski Security Extends Security Device
Support and Management Services to Fortinet**
Expanded expertise and certifications enable Kudelski
Security to offer manufacturer support and remote
device management for Fortinet security devices

**• AKAMAI recognized as a Leader for CDN in 2019
among the Top 8**
Company earns recognition for its comprehensive
solutions and services suite that includes web
performance, media delivery, cloud security and network
operations

• Data Integration Tools: Top 16 Vendors
By 2021, more than 80% of organizations will use more
than one data delivery style to execute their data
integration use cases

Dear Reader,

Here is your free copy of
Telecom COTS World
one of our six High Focus e-magazines

Your benefits:

• Selected subjects: HIGH FOCUS
• Easy reading
• TIME SAVER: includes many direct links
• UNIQUE: Dedicated Website for each
   e-magazine, remain visible for years
• FREE Distribution Worldwide
• FREE Subscription

Discover our 6 e-magazines
Click on the logos below

Semi Update World

Embedded Systems World

aiworld

Telecom COTS World
Broadband Broadcast IoT AI Convergence

IoT World

ATCA World

Visualizing Tech Giants' Billion-Dollar Acquisitions

November 1, 2019 | CBINSIGHTS

Our graphic shows every billion-dollar acquisition made by Facebook, Amazon, Microsoft, Google, and Apple, from Facebook's $22B purchase of WhatsApp to Amazon's $13.7B bet on Whole Foods.

Over the last 3 decades, the FAMGA tech giants — Facebook, Amazon, Microsoft, Google, and Apple — have collectively made over 750 acquisitions.

This cohort of acquirers has deep pockets. Some of their biggest checks have been written for high-profile companies such as career platform LinkedIn (acquired by Microsoft for $26.2B), chat app WhatsApp (Facebook, $22B), and video-sharing platform YouTube (Google, $1.7B).

But not all of these acquisitions have been unqualified successes. Microsoft ultimately wrote off its $7.2B Nokia deal, while Google paid $12.5B for Motorola Mobility (its largest acquisition to date) and sold the unit for a quarter of its acquisition price 2 years later.
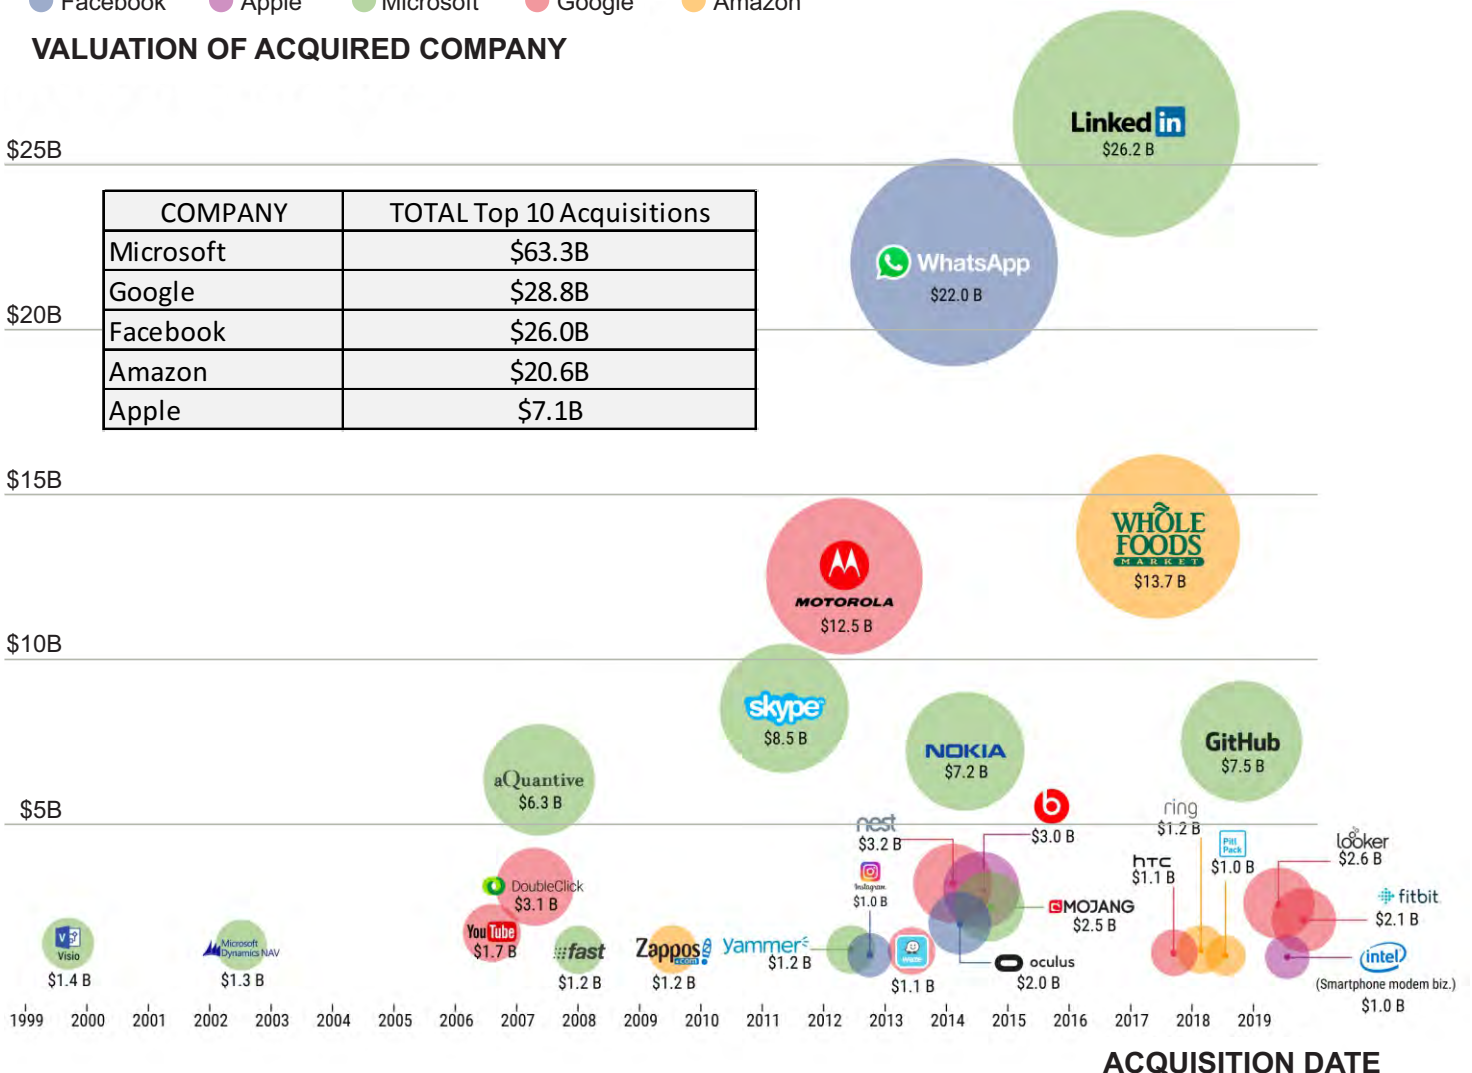
Using CB Insights M&A data, we made a visual timeline showing every $1B+ acquisition made by a FAMGA company.

TIMELINE OF TECH GIANTS' BILLION-DOLLAR ACQUISITIONS
Every $1B+ acquisition made by Facebook, Amazon, Microsoft, Google and Apple
Bubble size represents maximum valuation.

Legend: Facebook · Apple · Microsoft · Google · Amazon

VALUATION OF ACQUIRED COMPANY

COMPANY	TOTAL Top 10 Acquisitions
Microsoft	$63.3B
Google	$28.8B
Facebook	$26.0B
Amazon	$20.6B
Apple	$7.1B

ACQUISITION DATE

... to next page

Visualizing Tech Giants' Billion-Dollar Acquisitions

KEY TAKEAWAYS

Collectively, FAMGA has made 27 billion-dollar acquisitions.
Microsoft has made 10 $1B+ acquisitions, the most among these tech giants.
Apple has the least number of $1B+ acquisitions, with only its $3B acquisition of Beats Electronics and its recent $1B purchase of Intel's smartphone modem business.
Microsoft's $26.2B LinkedIn acquisition is the largest FAMGA buy, followed by Facebook's $22B WhatsApp deal and Amazon's $13.7B purchase of Whole Foods.
Microsoft's acquisition of Visio Corporation in 1999 was the first $1B+ acquisition by a FAMGA company.
The most recent $1B+ FAMGA acquisition was wearable manufacturer Fitbit, acquired by Google for $2.1B in November 2019.

Top 10 acquisitions by company

We broke down each of these companies' top 10 acquisitions by value. Read more about these deals in the posts linked below.

Facebook - Total: $26.0B
1. WhatsApp ($22B, 2014)
2. Oculus VR ($2B, 2014)
3. Instagram ($1B, 2012)
4. LiveRail ($500M, 2014)
5. Onavo ($200M, 2013)
6. Face.com ($100M, 2012)
7. RedKix ($100M, 2018)
8. Parse ($85M, 2013)
9. Pebbles Interfaces ($60M, 2015)
10. FriendFeed ($50M, 2009)
Read more about these acquisitions in our post, Infographic: Facebook's Biggest Acquisitions.

COMPANY	TOTAL Top 10 Acquisitions
Microsoft	$63.3B
Google	$28.8B
Facebook	$26.0B
Amazon	$20.6B
Apple	$7.1B

Microsoft Total is about the Total of the Next 4

Amazon - Total: $20.6B
1. Whole Foods ($13.7B, 2017)
2. Ring ($1.2B, 2018)
3. Zappos ($1.2B, 2009)
4. PillPack ($1B, 2018)
5. Twitch Interactive ($970M, 2014)
6. Kiva Systems ($775M, 2012)
7. Souq.com ($580M, 2017)
8. Quidsi ($500M, 2011)
9. Annapurna Labs ($370M, 2015)
10. LOVEFiLM International ($312M, 2011)
Read more about these acquisitions in our post, Infographic: Amazon's Biggest Acquisitions.

Microsoft - Total: $63.3B
1. LinkedIn ($26.2B, 2016)
2. Skype ($8.5B, 2011)
3. Github ($7.5B, 2018)
4. Nokia's Devices & Services Business ($7.2B, 2014)
5. aQuantive ($6.3B, 2007)
6. Mojang ($2.5B, 2014)
7. Visio Corporation ($1.4B, 1999)
8. Navision ($1.3B, 2002)
9. Yammer ($1.2B, 2012)
10. Fast Search & Transfer ($1.2B, 2008)
Read more about these acquisitions in our post, Infographic: Microsoft's Biggest Acquisitions.

Visualizing Tech Giants' Billion-Dollar Acquisitions

... from previous page

Google - Total: $28.8B
1. Motorola Mobility ($12.5B, 2012)
2. Nest Labs ($3.2B, 2014)
3. DoubleClick ($3.1B, 2007)
4. Looker ($2.6B, 2019)
5. Fitbit ($2.1B, 2019)
6. YouTube ($1.7B, 2006)
7. Waze ($1.1B, 2013)
8. HTC – Pixel Smartphone Division ($1.1B, 2017)
9. AdMob ($750M, 2009)
10. ITA Software ($700M, 2011)

Read more about these acquisitions in our post, Infographic: Google's Biggest Acquisitions.

COMPANY	TOTAL Top 10 Acquisitions
Microsoft	$63.3B
Google	$28.8B
Facebook	$26.0B
Amazon	$20.6B
Apple	$7.1B

Microsoft Total is about the Total of the Next 4

Apple - Total: $7.1B
1. Beats Electronics ($3B, 2014)
2. Intel — Smartphone Modem Business ($1B, 2019)
3. Dialog Semiconductor ($600M, 2018)
4. Anobit Technologies ($500M, 2011
5. Shazam ($400M, 2017)
6. NeXT Computer ($400M, 1996)
7. PrimeSense ($360M, 2013)
8. AuthenTec ($356M, 2012)
9. PA Semi ($278M, 2008) **see remark below**
10. Quattro Wireless ($275M, 2010)

Read more about these acquisitions in our post, Infographic: Apple's Biggest Acquisitions

This report was created with data from CB Insights' emerging technology insights platform, which offers clarity into emerging tech and new business strategies through tools like:
- Earnings Transcripts Search Engine & Analytics to get an information edge on competitors' & incumbents' strategies
- Patent Analytics to see where innovation is happening next
- Company Mosaic Scores to evaluate startup health, based on our National Science Foundation-backed algorithm
- Business Relationships to quickly see a company's competitors, partners, and more
- Market Sizing Tools to visualize market growth and spot the next big opportunity

If you aren't already a client, sign up for a free trial to learn more about our platform.

About PA Semi « What you may not know » By Daniel Dierickx, e2mos

PA Semi was a fabless semiconductor company founded in Santa Clara, California in 2003 by Dan Dobberpuhl (a guru in designing super CPUs) who was the lead designer for the DEC Alpha and StrongARM processors (HP acquired DEC). The company previously worked on processors like Itanium, Opteron and UltraSPARC. Apple Inc acquired PA Semi for $278 million in April 2008.

23 April 2008: Acquisition by Apple
On 23 April 2008, Apple announced that they had acquired PA Semi. While Apple's previous relationship with PA Semi would indicate that Apple could use their processors, PA Semi manufactures only Power ISA processors, which Apple does not currently use. At present, Apple only uses ARM and x86 processors.
On 11 June 2008, during the annual Worldwide Developer's Conference, Apple CEO Steve Jobs said that the acquisition was meant to add the talent of PA Semi's engineers to Apple's workforce and help them build custom chips for the iPod, iPhone, and other future mobile devices such as the iPad. PA Semi has said that they were willing to supply their PWRficient PA6T-1682M chip on an end-of-life basis, if the Power ISA license that PA Semi holds from IBM could be transferred to the acquiring company.

18 June 2007: 10 Months before the Acquisition by Apple
On 18 June 2007 Extreme Engineering Solutions published a PR introducing the Xcalibur1411 a Top Performance AdvancedTCA Blade for Telecom hosting eighteen PowerPC or nine dual-core PA Semi PA6T-1682 PWRficient processors.
In addition the Xcalibur1411 was a very Low Power Board.
In spite of the fact that Apple promised to support existing Customers (we all know that song) XES decided to stop the production of the Xcalibur1411.
X-ES (www.xes-inc.com) is a leading vendor of Embedded Computing Boards incl. COM Express, cPCI, VME, VPX & XMC/PMC

X-ES Extreme Engineering Solutions

Telecom COTS World - Nov-Dec 2019 - Page 5
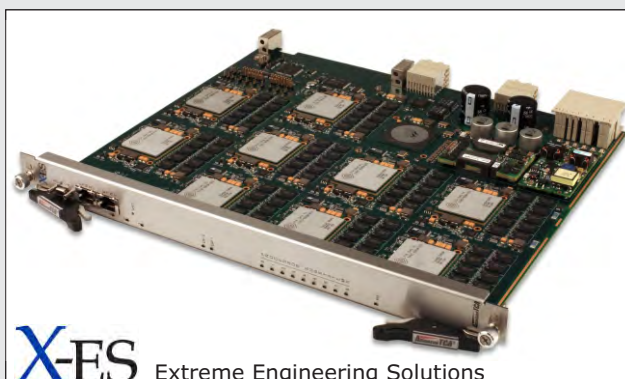
IASM - The Forrester Wave™:
Intelligent Application And Service Monitoring, Q2 2019

The 13 Providers That Matter Most And How They Stack Up
by Rich Lane with Laura Koetzle, Sandy Rogers, Julia Caldwell, and Diane Lynch

Today's Applications And Services Require Intelligent Monitoring

As complexity grows, I&O teams struggle to obtain full visibility into their environments and do troubleshooting. To meet rising customer expectations, operations leaders need new monitoring technologies that can provide a unified view of all components of a service, from application code to infrastructure. Forrester survey data shows that 51% of global infrastructure decision makers report they've already adopted, or are in the process of implementing, AI- and machine learning (ML)-enabled systems, with another 21% stating that they plan to obtain these technologies in the next 12 months.1 Vendors are expanding their offerings to provide a full range of monitoring capabilities across the stack and to integrate AI/ML for more efficient RCA with prescriptive and predictive guidance.

As a result of these trends, customers should look for IASM solutions that offer:

›› **Intelligent RCA and remediation.** Pinpointing the primary cause of an issue within a complex application technology stack can be frustrating and time-consuming. An intelligent monitoring solution that uses AI/ML can shorten response and remediation times by providing more accurate, prescriptive, and predictive guidance.

›› **Robust tools for measuring digital CX.** Real user monitoring (RUM), synthetic transaction monitoring, and transaction monitoring measure CX, tracking every step of a user's interactions, from device to application to ultimate business outcome.

›› **Ease of deployment and fast time-to-value.** With flexible monitoring solution architectures, deployment tools, and processes, dev and ops teams can push an agent in mere minutes to all the hosts they wish to monitor.2 And the less manual configuration an IASM solution requires, the faster an ops team can start using it to resolve problems.

Radisys Unveils Engage@Work:
A Communication and Digital Engagement Solution for Service Providers and Businesses

ʀadisys

One unique application, with multiple use cases – from customer and brand engagement to employee collaboration – securely and cost-effectively

HILLSBORO, OR, U.S. – October 9, 2019 – Radisys® Corporation, a global leader of open telecom solutions, today introduced Engage@Work – a communication and digital engagement solution that transforms how businesses engage with their customers and employees. Engage@Work offers interactive person-to-person and application-to-person public and private channel-based messaging and content delivery and person-to-person and multi-party high definition voice and video communication to facilitate collaboration and commerce. An industry-proven, web-scale solution, the underlying application platform that powers Engage@Work is currently deployed by a tier-one service provider in Asia where it supports tens of millions of monthly active users and over five million video calls every day.

News Highlights

• Communication and engagement with today's digital natives – customers and employees – must be mobile and take place in real-time, and it must be interactive, simple and contextual. This allows business communications to transform from just transactions to engagements, creating greater efficiencies while gaining crisp, actionable feedback. Traditionally, businesses have been forced to use multiple channels and applications to communicate with their audiences for brand engagement, customer care, community engagement, and employee collaboration. Engage@Work changes this paradigm: it offers a single extensible solution that brings together all communications and digital engagement under one application.

• An agile solution with multiple deployment options for service providers and system integrators to monetize digital communications, Engage@Work can be deployed as a turnkey, white-labeled and fully customizable application or as a plug-n-play software development kit (SDK) that can be integrated into pre-existing Enterprise applications. This enables service providers to leverage their networks and complementary services to reach their customers with more personalized services and flexible pricing and bundling while ensuring their customers greater control over user data.

• Engage@Work can be hosted in a service provider or enterprise network for total data privacy and user control.

• Engage@Work supports two-way interactive public and private channels to manage communications with large audiences, allowing businesses to send targeted and secure interactive messages with curated content to its customers and employees:

> For example, a business can broadcast a message to its user base one-to-many within a channel or to segments of that channel, and users can respond back to the business one-to-one.

> Public channels can support millions of users, while private channels can support secure communications among an employee base.

> Integration of bots within a channel can optimize workflow such as reduce the load on customer agents and increase efficiencies for customer service.

> Channels can be extended through connectors with third-party systems to expand the utility of the service.

> Radisys' Engage@Work is the only digital engagement solution that supports all of the following features:
  - Unlimited number of interactive channels to address all user demographics
  - Ease-of-use – audio and video conferencing capabilities on-the-go with no additional software to download
  - Ability to easily integrate widgets for productivity, commerce and more
  - A scalable mobile solution able to support service providers and large enterprises with up to 500 member groups
  - Support for all platforms – web, iOS, Android and feature phones and desktops and available on 2G-5G networks and Wi-Fi
  - Secure one-to-one interactions with consumers for a "no trolling" environment
  - Detailed analytics provide actionable data on follower metrics and response rates for channel administrators

"Consumers and employees want more personalized interactions with businesses where they shop and work. Businesses of all sizes need simple, secure real-time communication solutions that function across all devices without increasing the organization's operational workload," said Al Balasco, vice president, Communications Solutions, Radisys. "Engage@Work upgrades the digital engagement experience – communication, content, and commerce – for business and service providers alike, providing a seamless, feature-rich application that delivers a consistent and user-friendly digital experience for the mobile business community."

According to Gina Luk, Principal Analyst, Mobile Workforce Strategies at Strategy Analytics, "With the continued increase in mobile and distributed working, together with the growing need to collaborate across organizational boundaries, Radisys' Engage@Work not only transforms the way businesses communicate with their customer base, but also offers both reliability and security that safeguards enterprises against downtime."

Radisys ATCA T-100 Series

For more information, contact sales@radisys.com or visit www.radisys.com/engage/engageatwork

NEWS HIGHLIGHTS

● Wind River Cloud Platform combines a fully cloud native, Kubernetes and container-based architecture with the ability to manage a physically and geographically separated infrastructure for vRAN and core data center sites.

● Reducing service providers' operational burden and costs, the platform delivers single-pane-of-glass, zero-touch automated management of thousands of nodes.

● Cloud Platform is a commercially supported version of StarlingX and lends itself to demanding 5G use cases applicable across mission-critical industries.

KUBECON, SAN DIEGO, CA — Nov. 19, 2019 — Wind River®, a leader in delivering software for the intelligent edge, today announced Wind River Cloud Platform, a high-performance, production grade Kubernetes-based offering for managing edge cloud infrastructure. Optimized for the network edge, the platform addresses the service provider's complex challenges of deploying and managing a physically distributed, cloud native vRAN infrastructure.

"Telecommunication infrastructure has been evolving from vertically integrated, monolithic solutions to disaggregated technology based on virtualization and the cloud. Existing cloud infrastructure is unable to meet the new edge compute requirements that applications such as 5G, IoT, and MEC introduce," said Paul Miller, vice president of Telecommunications at Wind River. "Wind River Cloud Platform's revolutionary capabilities make 5G possible by solving the problem of deploying and managing distributed networks and is an ideal match to service providers' need for an operationally friendly, edge capable, multi-thousand node cloud. This new technology can become a de facto standard for production grade, distributed cloud infrastructure."

Cloud Platform leverages the OpenStack Foundation's open source project StarlingX to deliver the foundation for a geographically distributed managed solution able to simplify Day 1 and Day 2 operations by providing single-pane–of-glass (SPoG), zero-touch automated management of thousands of nodes, no matter their physical location. The platform provides orchestration of fully automated software updates and upgrades across a geo-distributed cloud, with rollback capabilities. Unlike enterprise-class IT platforms, Cloud Platform is an out-of-the-box open source solution designed specifically for edge compute that delivers ultra-low latency with deterministic performance. This is critical for high-availability, performance sensitive 5G virtualized radio infrastructure, as well as advanced applications such as autonomous vehicles.

"Wind River has been a long-standing contributor to open source projects. We are excited to have Wind River as a member of CNCF and we look forward to their contributions and collaboration to drive container technology to the edge," said Dan Kohn, executive director of Cloud Native Computing Foundation. "With Wind River Cloud Platform, Wind River is helping to further advance technologies such as Kubernetes at the edge."

"The scope of edge computing is vast and complex. StarlingX fulfills a real need for the right infrastructure to successfully realize cloud technology and services at the edge," said Jonathan Bryce, executive director at OpenStack Foundation. "With the introduction of Wind River Cloud Platform, we are delighted to see the rapid maturity and innovation of a commercial implementation of StarlingX."

Cloud Platform scales from a single compute node at the network edge to enable ultra-low cost deployments, up to thousands of nodes to meet the needs of high-value applications as they grow. Remote nodes can survive control plane disconnection and continue to operate and re-synchronize upon reconnection. All control functions can exist at all sites. Remote sites can be zero-touch enrolled and replicated across thousands of sites with fully automated deployment of known-good configurations.

Cloud Platform's common virtualization architecture can enable service providers to create a single approach to cloud infrastructure across far edge, near edge, regional, and core data center site elements within a planned 5G deployment topology. The distributed infrastructure can be managed as a single high-availability cloud, and do so with both containers and VMs co-existing seamlessly.

Wind River has for decades provided a backbone for global telecommunications infrastructure, with offerings used by all top telecommunications equipment manufacturers (TEMs). The company is a leader in the early 5G landscape, powering the majority of 5G RAN deployments. Now with Cloud Platform, Wind River can deliver, directly to service providers, one of the industry's most advanced cloud native distributed infrastructure solutions for 5G vRAN network deployment. Cloud Platform is a commercial implementation of the StarlingX open source project. StarlingX is a container-based cloud infrastructure software stack for edge implementations that demand ultra-low latency.

The rise of 5G connectivity also creates the potential to have a transformative impact across industries. Use cases in areas such as autonomous vehicles, drones, industrial automation/IIoT, and medical systems involve high stakes and demand more intensive requirements, specifically in regard to low latency, smaller footprint, and reliability. Wind River Cloud Platform can help companies deliver a distributed edge cloud solution for these demanding 5G use cases. With deterministic low latency required by edge applications and tools that make the distributed edge manageable, Cloud Platform provides a container-based infrastructure for edge implementations in scalable solutions ready for production. The unique presence Wind River holds across these industries ensures that service providers selecting edge vRAN 5G infrastructure have the ideal partner to enable the use cases driving 5G adoption.

**More information about Wind River Cloud Platform** is available at www.windriver.com/CloudPlatform.

Kontron ATX Motherboards "Designed by Fujitsu"

BIOS Updates support of the latest Intel® processors for better performance and longer lifetime

Augsburg, 11-Dec-2019 - Kontron, a leading global provider of IoT/Embedded Computing Technology (ECT), now offers a BIOS update to version R1.20.0 for its D3598-B and D3598-G motherboards, both "Designed by Fujitsu". This refresh now also supports the latest Intel® processors. The use of the latest CPUs increases the performance of the motherboards and extends their lifetime. In addition, more performance is offered at a lower price.
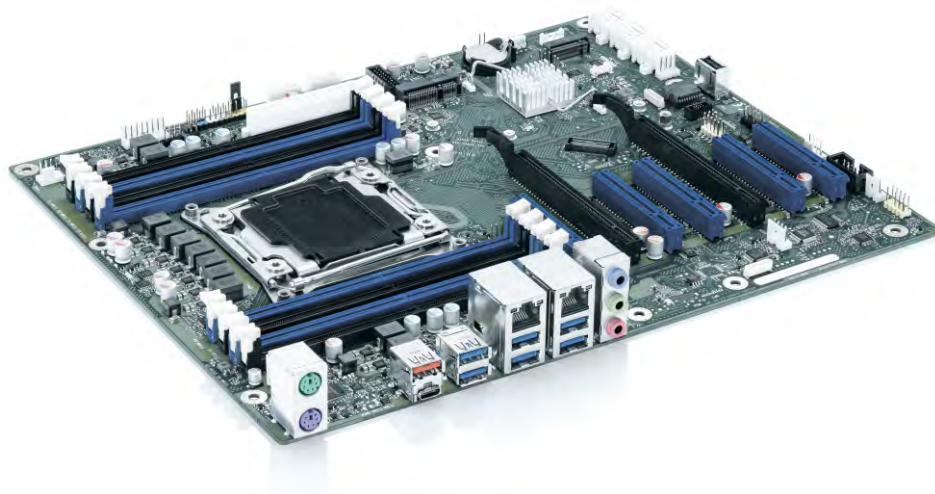
The Kontron D3598-B motherboard now supports the Intel® Xeon® processors of the W22xx series on LGA 2066 with the Intel® C422 workstation chipset. It meets semi-industrial requirements such as 24/7 operation at an ambient temperature of up to 50 degrees Celsius at full-load operation. Up to 512 GB R-DIMM/LR-DIMM (ECC) RAM can be accommodated on the motherboard and SLI certification for Nvidia Quadro exists for the use of a graphics controller.

The motherboard with the designation Kontron D3598-G can now also be operated with the tenth generation of Intel® Core™ processors on the Intel® X299 workstation chipset, such as the Intel® Core™ X-Series i9-109xxX. It has up to 256 GB UDIMM (non-ECC) memory.

Both motherboards offer 2xPCI Express® x16, Intel® VROC™ support, Dual LAN and Mini-PCI Express®, HW Watchdog and USB 3.1 Gen2 Type C. As with all Kontron Extended Lifecycle motherboards, there is an extended lifecycle of at least three years.

With their performance, features and robustness, the two ATX motherboards from Kontron, "Designed by Fujitsu", meet all requirements for demanding workstations or high performance computing.

For more information, please visit:

D3598-B ATX:

https://www.kontron.com/products/boards-and-standard-form-factors/motherboards/atx/d3598-b-atx.html

D3598-G ATX:

https://www.kontron.com/products/boards-and-standard-form-factors/motherboards/atx/d3598-g-atx.html

- HPE announced expanded intelligence and composability offerings by integrating HPE Primera storage HPE Synergy and HPE Composable Rack
- HPE Composable Rack supports VMware Cloud Foundation for hybrid cloud deployments
- HPE is enabling customers to deliver services on an intelligent cloud platform, offering the flexibility to support any application and service level with cloud-like agility

Offerings enable enterprises to support mixed workloads with greater speed, extreme data resiliency and ease of scale

SAN JOSE, Calif. – Nov. 4, 2019 – Hewlett Packard Enterprise (HPE) today announced combined intelligence and composability offerings by integrating its artificial intelligence (AI) and machine learning-driven HPE Primera storage platform with the composability in HPE Synergy and HPE Composable Rack1, helping customers rapidly deliver new apps and innovations to propel their businesses forward. This unique combination allows customers to deliver services on an intelligent cloud platform, offering the flexibility to support any application and service level agreement (SLA) with cloud-like agility, extreme resiliency, and seamless scalability.

Additionally, expanding on the recent introduction of HPE Synergy support for VMware Cloud Foundation (VCF), HPE Composable Rack, the HPE composable rack-scale solution, now supports VCF for hybrid cloud deployments. This new offering provides cloud management services to run enterprise applications in both private and public environments with a resilient and operationally efficient approach.

Enterprises today want to accelerate everything – their apps, data, and innovation. This requires the ability to more efficiently manage traditional bare-metal and virtualization applications, while also supporting containerized applications and cloud stacks with software-driven automation, and a fluid pool of resources that they can flexibly custom fit to the specific needs of applications. And, as organizations transition from a software-defined data center to an autonomous one, IT teams need to put the insights gathered about the infrastructure and workloads into action to ensure an environment that is always-fast and always-on.

"Customers want more than table stakes software-driven automation, and only HPE offers the combination of software- and intelligence-driven automation in a composable solution, transforming on-premises infrastructure to be more cloudlike," said Mat Mathews, VP and GM, Composable Networking, HPE. "And, as customers leverage VCF to provide integrated software-defined data center and cloud management services for their applications, it becomes critical to rely on an underlying platform that has matching characteristics. Only HPE's intelligent and composable infrastructure delivers the automation and flexibility needed for today's changing workload demands."

Composable Infrastructure

HPE is a pioneer in composability, initially introducing composable infrastructure in HPE Synergy, a bladed, modular architecture. With disaggregated compute, storage, and fabric, HPE composable infrastructure provides independent scaling of resources and on-demand provisioning to support any workload, and any SLA – all while delivering CapEx savings. Its software-defined automation improves IT efficiency and reduces OpEx costs, while its unified API allows developers to integrate infrastructure commands directly into development and deployment processes, as well as easily integrate with third-party applications and tools. Similarly, HPE Composable Rack provides disaggregated compute, storage and fabric resources in a rack form factor, supporting bare metal, containerization, virtualization, and private and hybrid clouds.

HPE advances the cloud experience through intelligence and composability

"Combined, HPE Synergy and our current storage solution, HPE 3PAR, have created a formidable capability for us to deliver industry leading services on composable and reliable infrastructure with a few clicks," said Christopher Horvat, VP of Information Technology, Skyview Networks. "We can virtualize all our audio in digital format without the traditional analog-to-digital latency and without a single physical backplane connection. Using this technology has made our organization more agile in delivering premier services to our long list of clients, so I look forward to the impact HPE Primera combined with HPE Synergy will have for mission-critical apps."

> **HPE Primera leverages the industry's most advanced AI platform for operations, HPE InfoSight**

Intelligent Infrastructure

HPE Primera leverages the industry's most advanced AI platform for operations, HPE InfoSight, to deliver significant breakthroughs – including 93% less time spent on managing storage2, the ability to predict and prevent issues, and accelerate application performance. HPE Primera lets organizations innovate without compromise, with the resiliency and performance businesses need combined with the agility of the cloud.

The integration of HPE Primera in the composable portfolio allows HPE Synergy and HPE Composable Rack customers to benefit from:

Agility and faster time-to-market

• Develop and deploy applications faster with infrastructure as code to automate and streamline resource management, including reduced time to deploy for Oracle databases by 89 percent3.
• Eliminate disruptions and wasted time reacting to problems with the AI-driven full-stack intelligence.
• Enable self-sufficiency and eliminate lag time with provisioning requests.

Extreme resiliency

• Through cloud-based machine learning and AI, HPE InfoSight in HPE ProLiant servers, HPE Synergy, and HPE Primera storage, predicts and prevents problems before they occur.
• Mission-critical reliability is built into HPE Primera, which is architected for high availability with its multi-node design, transparent business continuity with HPE Peer Persistence, and data replication—backed by the HPE Primera 100 percent data availability guarantee.
• HPE further protects systems against service disruption and risk, with its silicon root of trust that prevents system boots with compromised firmware; and Cyber Catalyst designation that provides enhanced terms and conditions on cyber insurance policies.

Seamless dynamic scaling

• HPE composable offerings allow organizations to compose disaggregated resources on demand and non-disruptively.
• The seamless auto-discover and -configuration capabilities enable incremental scale and integration of resources.
• The instant visibility into how resources are consumed by applications improves insights for capacity planning.

> **Through the unique combination of intelligence and composability, HPE provides an intelligent cloud foundation**

Through the unique combination of intelligence and composability, HPE provides an intelligent cloud foundation that unites the IT environment with a common software-defined and AI-driven operating model that can be optionally delivered as a Service via HPE GreenLake to align costs with consumption.

Pricing and Availability

HPE Synergy and HPE Composable Rack integration with HPE Primera; and support for VMware Cloud Foundation on HPE Composable Rack will be available in November 2019.

Additional Resources

• Visit this blog to read more about the new HPE composable capabilities announced today, and how they enhance the cloud experience on premises
• For a look at how composability and intelligence are accelerating the self-driving data center, read more here
• To learn more about how HPE Primera is advancing the cloud experience, visit this blog
• Full article from HPE: Click Here

Software-Defined and Cloud-Native Foundations for 5G Networks

Introduction

The fifth generation of mobile networks, commonly known as 5G, holds a lot of promise. Historically, 2G brought us mobile voice, while 3G introduced us to mobile data. 4G and LTE enabled usable mobile broadband services, and now 5G is supposed to unlock further value from our mobile networks with faster broadband data speeds and lower latency. 5G is expected to bring about a sea of change that will come as a result of new use cases and business models for consumers, businesses, and the industry. And it will lay the foundation for massive adoption of Internet of Things (IoT) everywhere.

5G and the applications it enables require a significant upgrade of carrier infrastructure beyond the radio access networks (RAN). Putting in 5G radios is just part of the overall effort. To support the diversity of applications, an explosion in device types, capacity and scale of traffic, communications service providers (CSPs) are turning to the adoption of cloud data center innovation both adapted and re-developed for carrier networks. Virtualization, software-defined infrastructure, cloud-native computing technologies in the form of network functions virtualization (NFV), software-defined networking (SDN), multi-access edge computing (MEC), containers and micro-services architectures will have significant roles to play in 5G infrastructure upgrades.

> Communications service providers (CSPs) are turning to the adoption of cloud data center innovation both adapted and re-developed for carrier networks.

5G - An Unstoppable Wave

Some might ask why cloud-native and software-defined technologies are needed within 5G networks. The reason lies in the impending scale of the mobile traffic that the 5G core (5GC) and the edge will have to process shortly. For instance, Ericsson's mobility research report estimates that mobile traffic is growing at a rate of around 40%, and Cisco's Visual Networking Index estimates that wireless and mobile devices will account for more than 63% of total IP traffic by 2021.

The wave has already started. AT&T recently rolled out 5G in 12 cities at the end of 2018, and will deploy service in 19 more markets in early 2019. Verizon has already rolled-out 5G fixed wireless service across four US cities and demonstrated live 5G at CES 2019. It plans to expand more aggressively in 2019. All carriers are laying the groundwork for 5G by increasing existing LTE capacities with solutions like LTE Advanced and LTE-LAA (Licensed Assisted Access), which provide 400Mbps to 1Gbps peak speeds. Moreover, with

Software-Defined and Cloud-Native Foundations for 5G Networks

5G data-only devices starting to become available, operational 5G networks are just around the corner.

In addition, we expect that module-based 5G devices for IoT will come online around 2020. We already have significant IoT traffic on today's mobile networks, from surveillance cameras to agricultural sensors, to building HVAC monitors. Research firm Statista predicts that 12.86 billion IoT sensors and devices will be in use in the consumer segment by 2020, growing at 35% CAGR from 2017. Whether LTE or 5G, IoT devices will make up a good number of mobile end-points.

What all this means at the end of the day is increased traffic, an increased number of devices and higher demands on the mobile edge, transport and core infrastructure. This all leads to more challenges for global telcos.

> What all this means at the end of the day is increased traffic, an increased number of devices and higher demands on the mobile edge, transport and core infrastructure.

New 5G-Enabled Use Cases

So what can 5G enable that 4G can't handle? The 3GPP has defined a large number of different 5G use cases, categorized into a few different groups including massive IoT, critical communications, enhanced mobile broadband (eMBB), and network operation. Also, the Next Generation Mobile Networks (NGMN) Alliance, in its February 2015 whitepaper, described 25 use cases grouped into eight families: broadband access in dense areas, broadband access everywhere, higher user mobility, massive IoT, extreme real-time communications, lifeline communication, ultra-reliable communications,and broadcast-like services.

All these new use cases depend on the lower latency, higher-bandwidth and increased flexibility that 5G promises. The mobile consumer has been bombarded with images of new applications like augmented reality/ virtual reality (AR/VR), connected and autonomous vehicles, telemedicine, and smart cities. So consumers can't wait to get their hands on these new applications, but there's work on the backend to make all this happen. For instance, to support realistic AR/VR, the network needs to provide sub-10ms latency and high bandwidth. Moreover, remote surgery in telemedicine requires extremely low latency to be successful. Meanwhile, premium HD, 360° and 4K video require both high-bandwidth and the support for edge caches.

5G and Network Slicing

To support these different use cases, 5G infrastructure provides for a capability known as network slicing. This is the ability to enable multiple separate networks on top of a common shared infrastructure. From a carrier perspective, a network slice is a virtualized end-to-end logical network on top of a physical infrastructure that provides a specific quality of service that is negotiated during the provisioning of the network.

This network slice might use dedicated physical resources, or perhaps shared resources, from the base station of the radio access network all the way through the transport layer and into the core.

One of the use cases for network slicing is to support mobile virtual network operators (MVNOs). With virtual network slices, the extended networks of multiple virtual carriers or even enterprises can stretch beyond enterprise boundaries into the mobile sphere as well, running on top of the shared physical infrastructure.

Other examples of network slices include a high bandwidth slice for movie streaming or an ultra-low latency slice for telemedicine. Another example would be an ultra-reliable slice for autonomous driving or assisted driving.

Network slicing can facilitate new business models for carriers to sell or share parts of the network in a secure and isolated manner to different companies. This is similar to AWS or Microsoft Azure selling cloud services as shared resource slices of computing, networking, and storage to various enterprises.

> Network slicing can facilitate new business models for carriers to sell or share parts of the network in a secure and isolated manner to different companies.

Key Technologies in the Cloud-Native World

Speaking of cloud technologies, to bring innovative new 5G applications and network slicing to market, carriers have realized that they need to turn to new technologies. Carriers and industry organizations like 3GPP and NGMN have recognized that their best bet lies in adopting technologies pioneered within cloud data centers. Only through adopting these technologies can carriers provide the scale and performance required by 5G applications.

SDN

SDN is a movement that started in the early 2010s. It involves the separation of the control plane of network devices from the data plane, allowing a centralized approach for networking control that provides simplification and global optimization for the routing and switching of network packets. SDN also advocates open APIs and a programmatic approach to networking.

SDN is an essential element of 5G that enables fast service provisioning (and de-provisioning) as well as the optimal use of the underlying transport. Through its support for a programmable network, SDN is critical in virtualization of the underlying network.

SDN architectures can also be implemented to ensure that end-to-end paths are provisioned efficiently to maximize the transport from the mobile user equipment (UE) to the data center or edge services that they are likely to consume.

Software-Defined and Cloud-Native Foundations for 5G Networks

NFV

SDN itself will facilitate the coordination and virtualization of the underlying network, but a 5G network also requires network functions that support critical elements in each network slice. Specifically, network slicing will require flexible network services per slice that perform the kind of network functions that are found today in the SGi/Gi-LAN, such as firewalling and other security functions; caching and acceleration; and any metering or policy elements.

NFV is a complementary movement to SDN, leveraging virtualization to take proprietary physical hardware that is non-portable and hard to manage, and convert the network functions into virtualized software-only versions. These software virtual network functions (VNFs) can be quickly moved around as needed, and scaled up or down dynamically.

NFV architectures support the ability to provision network slicing services flexibly, and this will likely take the form of data centers spread across the network from core to edge to reduce the latency of these services and for load distribution.

> The combination of NFV and SDN is critical in achieving flexible network topology and the realization of 5G targets.

Even in core 5G networks that don't need network slicing, it is still essential to have the flexibility to create network services. Virtualized mobile cores (vEPC) enable flexible network management while maximizing its availability based on virtualization aspects.

The combination of NFV and SDN is critical in achieving flexible network topology and the realization of 5G targets such as 1,000-times higher system capacity; 100-times increase in data rates (10-Gb/s speeds); connectivity enablement for 100-times more devices; latency reduced to 1 millisecond from 5 ms; and energy savings.

Containers

SDN and NFV are well-understood and already deployed in today's 4G LTE networks. However, to enable the next level of flexibility in mobile networks, operators might have to turn to more recent innovations in the data center world — containers and cloud-native capabilities.

One of the most fundamental changes to occur in cloud platform the last few years has been the ascendency of containers and their associated microservices. While not necessarily a new concept, containers and container solutions like Docker, have made it much simpler for developers to package their applications in a way that makes it possible for them to be deployed, both on-premises and in the cloud, at will.
By providing a higher level of abstraction, containers eliminate the need for developers to navigate multiple types and classes of virtual machines.

Containers allow applications to be separated from the underlying infrastructure, making them more portable. They also provide speed and agility to allow applications to be tested and deployed quickly. Further, containers provide lower overhead than virtual machines while still providing isolation from infrastructure and other tenants on the same hardware and operating system.

Microservice Architecture

Historically, applications were written as a monolithic entity. The more complex the application, the larger the deployable binary. Microservices architecture is an approach that breaks a monolithic application down into a collaborating collection of components, called microservices. For instance, in an application that involves a database, there might be a microservice that is responsible for search and another responsible for backup. The goal is to simplify the application into many components, usually running on containers, each of which can be independently deployed, upgraded, and patched. The thought behind this architecture is that each component tier can be scaled accordingly to accommodate scale, and each component can be upgraded independently to facilitate agility.

To orchestrate the bring up of multiple components running on containers, there are a few popular orchestration systems, with Kubernetes from Google being the dominant solution in the cloud community today. Kubernetes can orchestrate the instantiation of hundreds or more separate components necessary to instantiate an application.

> Each microservice component tier can be scaled accordingly to accommodate scale, and each component can be upgraded independently to facilitate agility.

Applications of Cloud Frameworks in 5GC

To tie all these elements together, we need to examine the 5G system architecture as described by the 3GPP's Technical Specification (TS) 23.501 working group. Historically, in 2G, 3G and 4G networks, the different network functions that come together (e.g., HSS, MME, etc) are integrated as point-to-point links. This point-topoint model has served the mobile industry well in the past.

However, going forward, for 5G, the 3GPP has recognized that there is a need for increased service agility. As such the 3GPP working group has designated a services-based architecture (SBA) that mirrors cloudnative architectures and is believed to scale better.

> Just like containers and micro-services architectures did for cloud computing, SBA will facilitate 5G network functionality becoming more granular and decoupled.

Software-Defined and Cloud-Native Foundations
for 5G Networks

SBA Enables Scalable Network Services for 5G

Just like containers and micro-services architectures did for cloud computing, SBA will facilitate 5G network functionality becoming more granular and decoupled. This should allow for increased automation and an agile operational process (not unlike the continuous integration/continuous deployment or CI/CD movement in the cloud computing world). The result should be an overall reduction in deployment times, better operational efficiencies and improved resilience.

Within the context of SBA, a service is an atomic entity (like a microservice) that serves a specific function. A service can be updated independently and deployed as needed. In release 15, 3GPP introduced the concept of Network Functions Services (NF Services) as part of SBA. In this first release, SBA only applies to controlplane functions but will evolve beyond in release 16.

Under SBA, the 3GPP has defined common supporting infrastructure to make service-based architecture a reality. These capabilities include service registration, which provides a list of available services, their status and how to reach them. They also include service authorization and authentication which controls whether services can contact each other and service discovery, which provides for appropriate selection of the right services to use, and can facilitate load-balancing.

Also, 3GPP has picked common protocols (in release 15) reminiscent of cloud architectures for SBA. These currently include an overall RESTful framework utilizing HTTP/2 for application layer communication, TCP-based transport and JSON for serialization.

As depicted in the SBA diagram from the 3GPP above, the major components of the 5G core are listed as follows:

• Access and Mobility Management Function (AMF): Manages all UE related functions, especially access control and mobility.

• Session Management Function (SMF): Session establishment, modification, and release, as well as other control plane functions.

• User Plane Function (UPF): Provides functions specific to U-plane processing, act like the S and P Gateway in 4G. UPFs can be deployed in different locations to perform different functions.

• Policy Control Function (PCF): Equivalent to PCRF in 4G. Dictates the policy that governs overall network behavior.

• Unified Data Management (UDM): Stores subscriber data and profiles.

• NF Repository Function (NRF): Provides a service registry and service discovery for Nfs.

• Network Exposure Function (NEF): API gateway that allows external services or users to integrate with the mobile network and provision or deprovision services as well as push application policy. Works like the Service Capability Exposure Function (SCEF) in 4G.

• Authentication Server Function (AUSF): Supports authentication services.

These and other services identified above are meant to recreate under the SBA framework the components that were available within the pre-5G mobile core.
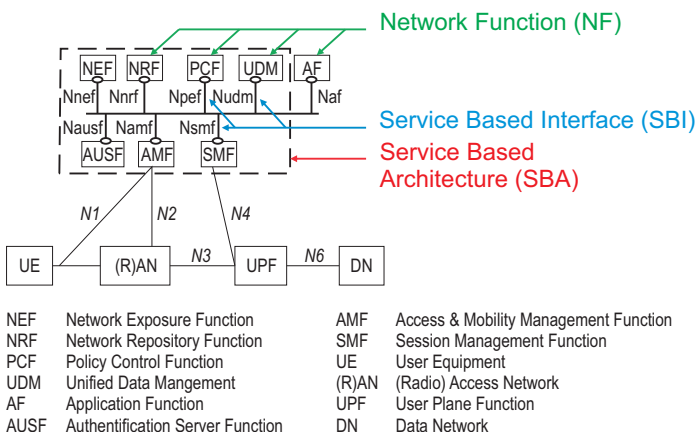
Network Slicing, SDN, NFV, and SBA - Tying it Together

Today, outside the mobile core, SDN, and NFV are used in conjunction to provide network functions within multitenant frameworks. By using SDN to provide network virtualization and service function chaining, a specific tenant can be served by a subset of available VNFs in an NFV deployment by merely associating a specific chain of VNFs for flows belonging to that tenant. This framework works regardless of whether these NFs are VNFs or CNFs (cloud-native network functions packaged as containers).

> The good news is that numerous technologies can be borrowed from the cloud world, including Kubernetes, service meshes, message bus frameworks, monitoring, and telemetry, and distributed key-value stores to help make SBA a reality.

Service Based Architecture

NEF	Network Exposure Function	AMF	Access & Mobility Management Function
NRF	Network Repository Function	SMF	Session Management Function
PCF	Policy Control Function	UE	User Equipment
UDM	Unified Data Mangement	(R)AN	(Radio) Access Network
AF	Application Function	UPF	User Plane Function
AUSF	Authentification Server Function	DN	Data Network

Figure: 5G Service Based Architecture (source: 3GPP)

Software-Defined and Cloud-Native Foundations for 5G Networks

Taking the same approach within the SBA framework, 5G vendors have set up an early demonstration of a similar setup on the 5G control plane. For instance, a mobile operator can request a network slice for surveillance video feeds that consist of a next-generation firewall to protect these IoT devices from malware, followed by some video optimization function. At the same time, this network slice can be allocated resources to ensure that it has sufficient bandwidth to deliver the videos while not necessarily needing the best latency.

> Just as AWS, Microsoft Azure and Google Cloud Platform use containers to power some of the world's most successful web services, the vision is to use the same technologies to enable even more massive scale at incredible efficiencies on 5G networks.

Under the SBA framework, and in concert with SDN and NFV, we can envision an example with the instantiation of two different NF services, such as AMF servicefirewall and AMF-service-video-opt solely for this network slice through the use of VNFs or CNFs under NFV. We would then orchestrate the underlying network via SDN to ensure the traffic from this slice, as it comes through the RAN via the transport network, is routed to these AMF instances running within our data center, for appropriate processing.

While this is going on it is possible to leverage the same technologies to simultaneously service a telemedicine network slice that requires different NFs and extremely low latency using the same shared physical resources.

Also, when one of the slices is longer needed, the services would be de-provisioned dynamically, and the resources returned to the shared pool.

Indeed, there is a lot more work needed before all this becomes a reality. Resource management, admission control on whether a network slice can be instantiated to guarantee the appropriate SLAs, what to do when a node in the SBA fails, and SLAs are put at risk. How to best orchestrate the AMFs and enable efficient service discovery etc. are all questions that must be resolved. The good news is that numerous technologies can be borrowed from the cloud world, including Kubernetes, service meshes, message bus frameworks, monitoring, and telemetry, and distributed key-value stores to help make SBA a reality. We've only seen the tip of the cloudnative iceberg within the 5G ecosystem.

Conclusion: Looking Towards Release 16 and Beyond

While we are early in these demonstrations, the hope is that when Release 16 is frozen in March 2020 (delayed from December 2019), the full concept of SBA has been fleshed out and that we are using a microservicesbased container infrastructure (perhaps orchestrated by Kubernetes) to power multiple 5G network slice at scale. Just as AWS, Microsoft Azure and Google Cloud Platform use containers to power some of the world's most successful web services, the vision is to use the same technologies to enable even more massive scale at incredible efficiencies on 5G networks.

About SDxCentral

SDxCentral's mission is to arm technology professionals with the information and context they need to understand the rapidly evolving IT infrastructure market, enabling them to lead their organizations into the technology-driven future. We deliver independent content that reaches our audience where they are, cuts through the marketing fluff, and combines veteran voices with fresh perspectives. MORE: please visit https://www.sdxcentral.com/

Enea Unveils Complete 5G Data Management Suite

Cloud-native 5G data suite to support automated network processes and data monetization

STOCKHOLM, Sweden – December 12, 2019 - Enea® (Nasdaq Stockholm: ENEA) today unveiled its end-to-end 5G Data Management suite, enabling interoperability between network functions in a software-defined 5G core. Following the successful acquisitions of Openwave Mobility and the Telco Solutions Convergence Creators business unit from Atos/Siemens, Enea has integrated and updated its cloud-native solutions to give operators a best of breed option for management of data in 5G networks.

The Enea 5G Data Management portfolio gives operators a unified view of their data and allows all applications for network and business operations to securely store and access data through open APIs. The Enea 5G Data Management suite includes 5G core network capabilities spanning the common or network data layer (CDL/NDL), the user plane and the control plane. It includes functions such as Unified Data Management (UDM), Unified Data Repository (UDR), Authentications Server Function (AUSF), Policy Control Function (PCF), and Equipment Identity Register (EIR).

One of the key components in the Enea 5G Data Management portfolio is Stratum, a Common Data Layer solution, which is already being used by several of the world's largest mobile operators. By collapsing all system data silos into a single common data layer, Stratum allows all applications to draw from the same pool of data. This effectively ends the major frustration of vendor lock-in experienced in 3G and 4G networks, based on data being siloed in different network functions.

"A cloud-native data management environment, such as the Enea 5G Data Management suite, will be an integral part of the new 5G networks which call for a highly automated, open architecture" said Gorkem Yigit, Senior Analyst at Analysys Mason. "Operators will need data management solutions that can not only meet the low latency and high-performance requirements of 5G but also store and expose critical data across all 5G services through open interfaces and support the deployment of multi-vendor and multi-generation interworking scenarios."

The availability of the Enea 5G Data Management suite supports many of the mobile operator community's expectations for 5G. The suite will deliver the foundation for the drive towards greater automation, as all applications for network and business operations will have the decision-making data they need. Through two-way open interfaces, the 5G Data Management suite enhances interoperability between vendors and delivers greater control for operators. It also provides a means for greater visibility and visualization of the network, which will be critical for network slicing to deliver improved monetization from enterprise and Internet of Things use cases.
 Last, but not least, flexible interworking with 4G and other networks on subscriber data and user plane level allows smooth integration and migration.

"Key 5G applications require robust data models and structures as a foundation for effective network automation and data monetization", commented Enea CEO, Jan Häglund. "Operators can realize the full potential of 5G only by leveraging best-in-class data management".

About Enea

Enea is a global software company with a product portfolio that spans the network from the edge to the core. The company provides critical software components that help realize operators' ambitions of a cloud-native, software-defined network.

More information about Enea 5G Data Management: https://www.enea.com/products/5g-data-management/

Kudelski Security Extends Security Device Support and Management Services to Fortinet

Expanded expertise and certifications enable Kudelski Security to offer manufacturer support and remote device management for Fortinet security devices

Cheseaux-sur-Lausanne, Switzerland and Phoenix (AZ), USA – Kudelski Security, the cybersecurity division within the Kudelski Group (SIX:KUD.S), today announced the expansion of its Fortinet Security Device Management Service by adding optional device support services, for which the global security firm provides from its Cyber Fusion Center in the United States.

Kudelski Security provides manufacturer support services for a broad range of vendors including F5, Juniper and Pulse Secure, and has over 120 clients across the United States. As part of Kudelski's 24x7x365 service, clients receive a premium service for quicker escalation and resolution, including finding, escalating, and tracking bug fixes. Key client satisfaction metrics include 99 percent renewal rate and maintaining a 100 percent service level attainment over the past 3 years.

"Fortinet has a powerful technology portfolio that features impressive performance and broad integration with products from other manufacturers," said Rich Fennessy, chief executive officer, Kudelski Security. "By adding manufacturer support for product issues we're able to extend greater value and service our collective clients, including expanded service ranges and faster response times. This frees our technology partners to stay focused on new features and major enhancements."

Kudelski Security has strongly grown its device support and device management services, which are part of its larger portfolio of advanced Managed Security Services (MSS), including Endpoint Detection and Response, Threat Hunting, and Attacker Deception. Kudelski Security was also named as a leader in the latest "Forrester Wave™: Emerging Managed Security Services Providers (MSSP), Q3 2018" and as a key player in Gartner's "Europe Context: Magic Quadrant for Managed Security Services, Worldwide" report.

"The service levels and advanced security services we are providing clients clearly differentiate our capabilities against legacy MSSPs," continued Fennessy. "This is the fastest growing area of our business and we continue to expand coverage, adding new clients around the world."

For more information about Kudelski Security's robust MSS offering visit https://www.kudelskisecurity.com/services/managed-security .

About Kudelski Security

Kudelski Security is the premier advisor and cybersecurity innovator for today's most security-conscious organizations. Our long-term approach to client partnerships enables us to continuously evaluate their security posture to recommend solutions that reduce business risk, maintain compliance and increase overall security effectiveness. With clients that include Fortune 500 enterprises and government organizations in Europe and across the United States, we address the most complex environments through an unparalleled set of solution capabilities including consulting, technology, managed security services and custom innovation.

For more information, visit www.kudelskisecurity.com .

AKAMAI RECOGNIZED AS A LEADER IN THE IDC MARKETSCAPE: WORLDWIDE COMMERCIAL CDN 2019 VENDOR ASSESSMENT

Company earns recognition for its comprehensive solutions and services suite that includes web performance, media delivery, cloud security and network operations

Cambridge, MA | August 29, 2019

Akamai (NASDAQ: AKAM), the intelligent edge platform for securing and delivering digital experiences, today announced that it has been positioned in the Leaders category in the IDC MarketScape: Worldwide Commercial CDN 2019 Vendor Assessment. As a Leader in the report, Akamai is recognized for its market reach, capabilities and innovation in the areas of web performance, media delivery, cloud security and network operations across its highly distributed global CDN infrastructure.

"Major CDN (Content Delivery Network) providers are refocusing their effort on the edge to address new demands for real-time video streaming, security and IoT," said Ghassan Abdo, research VP, IDC. "Akamai is a leader in delivering edge services, provides its content delivery services to a vast number of global enterprises and has comprehensive solutions and services offerings for web performance, media delivery, cloud security, and network operator solutions."

To compile the report, IDC identified and examined eight providers by scale and scope that provide CDN services across the globe. Based on structured discussions, surveys and interviews with market leaders, participants and end-users, the study evaluates these CDN providers in terms of their ability to offer services and strategies to grow and innovate in the marketplace.

IDC MarketScape Worldwide Commercial CDN 2019

As a result of this evaluation, Akamai was named a Leader in the CDN market assessment, recognized for its established position and largest footprint in the market, as well as several differentiators. Specifically, the report highlights the following strengths:

- **Capacity, performance and reliability assurance:** Akamai's network capacity is estimated to be around 90-120 TBps. Akamai delivers average daily web traffic of more than 40 TBps, and peak traffic exceeded 83 TBps in March 2019.
- **Security:** Akamai has had a robust edge security offering for quite some time and clearly aims to further expand its cloud security services offerings. Akamai has extensive threat intelligence and insights, interacting with 130TB of data, over 1 billion devices and 100+ million IP addresses daily.
- **Media Delivery:** Akamai has a comprehensive media delivery suite of services, which it continues to enhance in key areas such as live event streaming delivery, managed services, performance and latency improvements.
- **Innovation:** Akamai continues to expand its services and capabilities beyond CDN and leverages innovative technologies such as artificial intelligence and machine learning to address new areas like IoT, DevOps and blockchain.
- **Services teams:** Underpinned by highly qualified services teams and its reach into many verticals, Akamai is very well capable in guiding its clients on how to deploy and leverage its CDN and edge services to reach its full potential and proven business benefits.

"For more than two decades, Akamai has helped customers deliver content that engages audiences, drives business and provides flawless online experiences on any device, anywhere and anytime," said Dr. Tom Leighton, CEO of Akamai. "From web performance and media delivery, to cloud security and network operator solutions, this is a dynamic industry. Being named a CDN Leader by IDC validates our ongoing commitment and tireless innovation to provide market-leading solutions that help our customers securely conduct business on the internet."

To download a copy of the IDC MarketScape: Worldwide Commercial CDN 2019 Vendor Assessment CLICK HERE
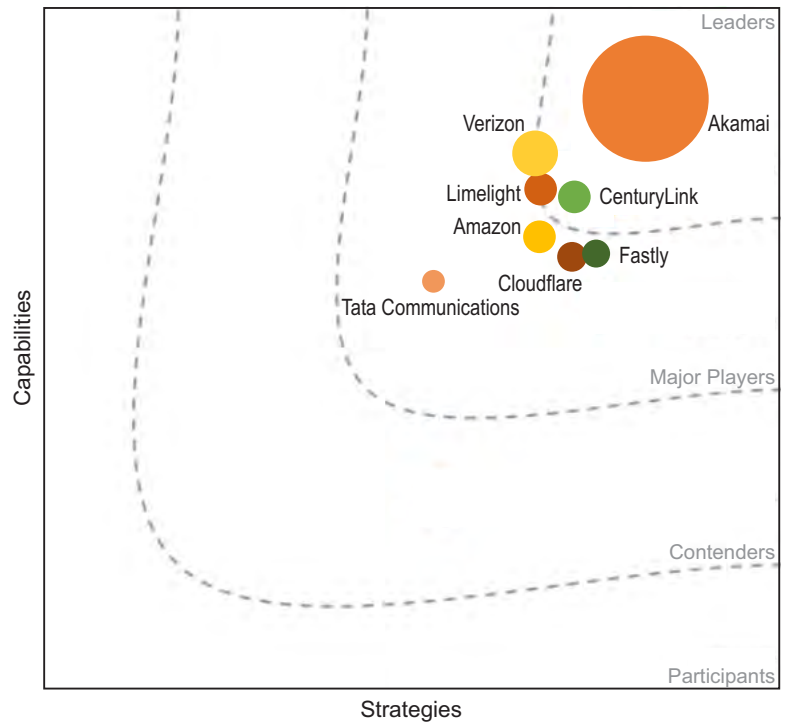
Gartner Magic Quadrant for Data Integration Tools

The data integration tool market is resurging as new requirements for hybrid/intercloud integration, active metadata and augmented data management force a rethink of existing practices. This assessment of 16 vendors will help data and analytics leaders make the best choice for their organization.

Strategic Planning Assumptions

● **By 2021**, more than 80% of organizations will use more than one data delivery style to execute their data integration use cases.

● **By 2022**, organizations utilizing active metadata to dynamically connect, optimize and automate data integration processes will reduce time to data delivery by 30%.

● **By 2022**, manual data integration tasks (including recognition of performance and optimization issues across multiple environments) will be reduced by 45% through the addition of ML and automated service-level management.

● **By 2023**, improved location-agnostic semantics in data integration tools will reduce design, deployment and administrative costs by 40%.

Market Definition/ Description

The discipline of data integration comprises the architectural techniques, practices and tools that ingest, transform, combine and provision data across the spectrum of data types. This integration takes place in the enterprise and beyond — across partners as well as third-party data sources and use cases — to meet the data consumption requirements of all applications and business processes. This is inclusive of any technology that supports data integration requirements regardless of current market nomenclature (e.g., data ingestion, data transformation, data replication, messaging, data synchronization, data virtualization, stream data integration and many more).

The market for data integration tools consists of vendors that offer software products to enable the construction and implementation of data access and delivery infrastructure for a variety of integration use-case scenarios.

Example integration usage scenarios include:

• Data integration and delivery for optimized analytics
• Sourcing and delivery of master data in support of master data management (MDM)
• Data consistency between operational applications
• Interenterprise data acquisition and sharing
• Data services orchestration
• Data migration and consolidation
• Support for data governance and management of data assets

Some examples of new and differentiating functionality or characteristics include:

• Interoperating with application integration technology in a single solution architecture
• Enabling data services for use in broader architecture approaches
• Supporting the delivery of data to, and the access of data from, a wide variety of data stores, repositories and data management tiers in application deployments
• Nonrelational DBMS integration
• IoT/OT data convergence
• Self-service data integration