PDF Tricks For The Linux Command-Line

A Quick Look At New Age Trends In AI

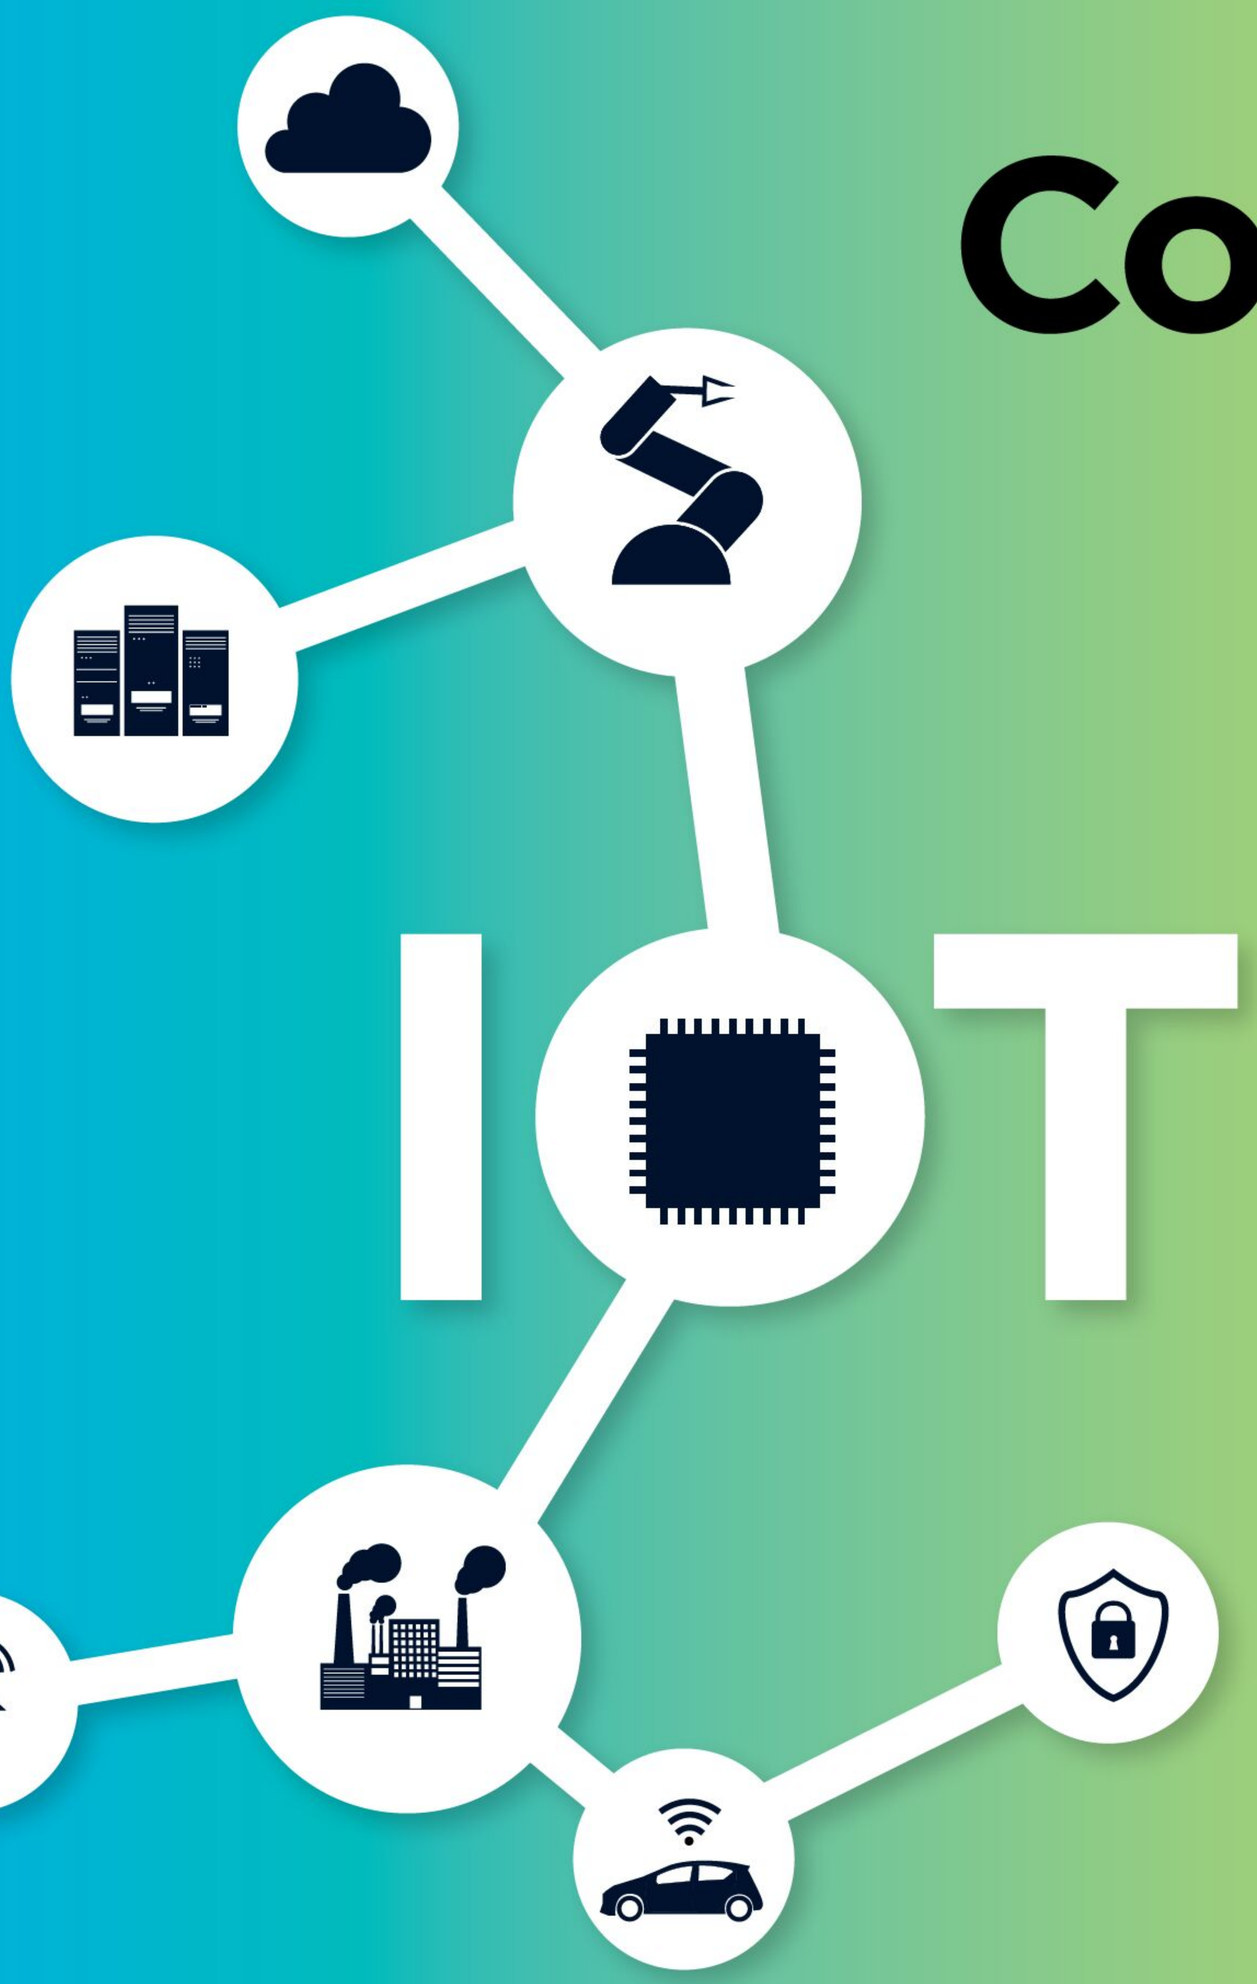
₹120
ISSN-2456-4885

OpenSource ForYou

Volume: 12 | Issue: 01 | Pages: 100 | November 2023

THE COMPLETE MAGAZINE ON OPEN SOURCE

An EFY GROUP Publication



Living With Connected Devices

Neuromorphic Computing: The Future Of IoT

Leveraging The Power Of CircuitPython For IoT Devices

Generative AI In IoT: Moving Towards Transformative Applications

Open Source India 2023: A Peek Into All That Happened

Simplifying The World Of Distributed Ledger Technology

www.opensourceforu.com

Wanna Be Your Own Boss?

DO OPEN SOURCE.



Demand for Open Source is sky rocketing. Be it for managing IT infrastructure or development of software--Open Source solutions are what customers are seeking.

All you need to do is develop expertise in an Open Source stack, and then build a team around it!

And, Open Source For You can be your friend and a guide through this journey.

TO READ OUR PRINT EDITION **Visit:** https://subscribe.efyindia.com

TO READ OUR EZINE EDITION **Visit:** https://ezine.lfymag.com



WORLD'S LEADING PUBLICATION ON OPEN SOURCE

Looking for marketing solutions to engage with cutting edge techies?
Contact us at growmybiz@efy.in OR call us at +91-9811155335.

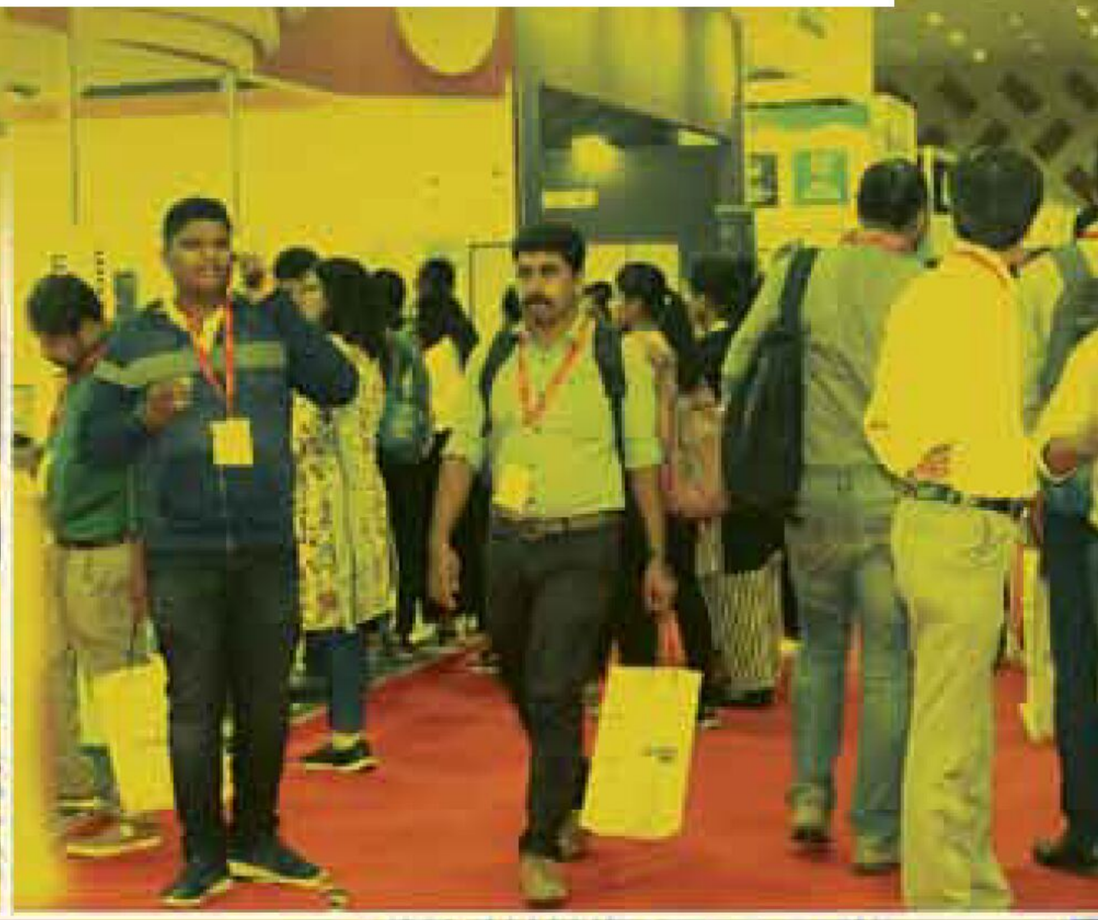
India's #1 Event For **R&D and New Product Development** Of Electronics

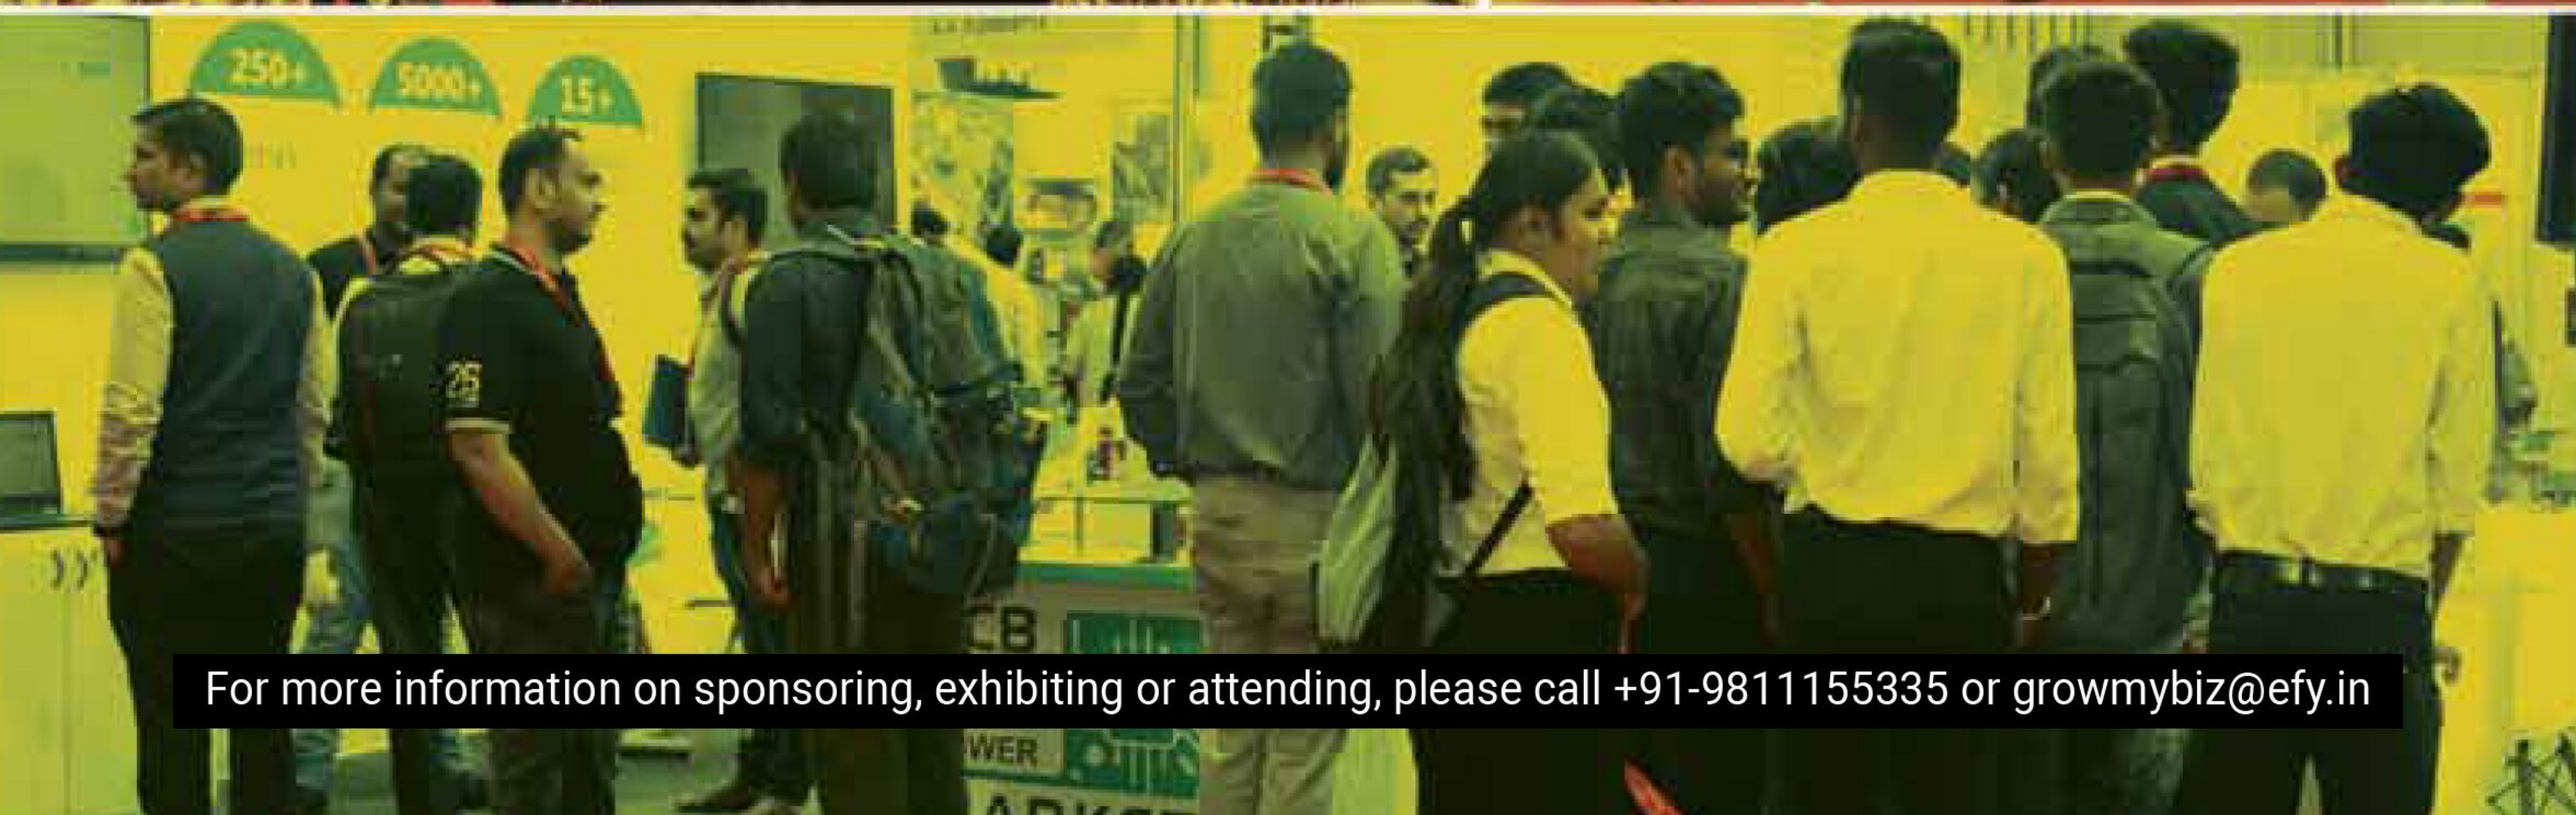
9TH EDITION INDIA ELECTRONICS WEEK 2024

31st Jan - 2nd Feb 2024

KTPO Expo Center | Whitefield | Bengaluru | India







For more information on sponsoring, exhibiting or attending, please call +91-9811155335 or growmybiz@efy.in

CONTENTS

NOVEMBER 2023 | ISSN-2456-4885

FOR U & ME

16 PDF Tricks for the Linux Command-Line

18 Fintech: The Digital Evolution of Finance

20 A Quick Look at New Age Trends in AI

32 The Benefits and Drawbacks of OSS

FOCUS

55 Generative AI: The Things It Can Help With

58 IoTivity: Using Open Source to Make Devices Talk to Each Other

65 Smart Healthcare: How IoT is Reshaping the Industry

69 How Open Source is Helping IoT Evolve

73 Leveraging the Power of CircuitPython for IoT Devices

75 Choose the Right Cloud Platform for Implementing IoT Protocols

DEVELOPERS

89 Revolutionising Land Management with Blockchain Technology

92 R Series: Caret Package

ADMIN

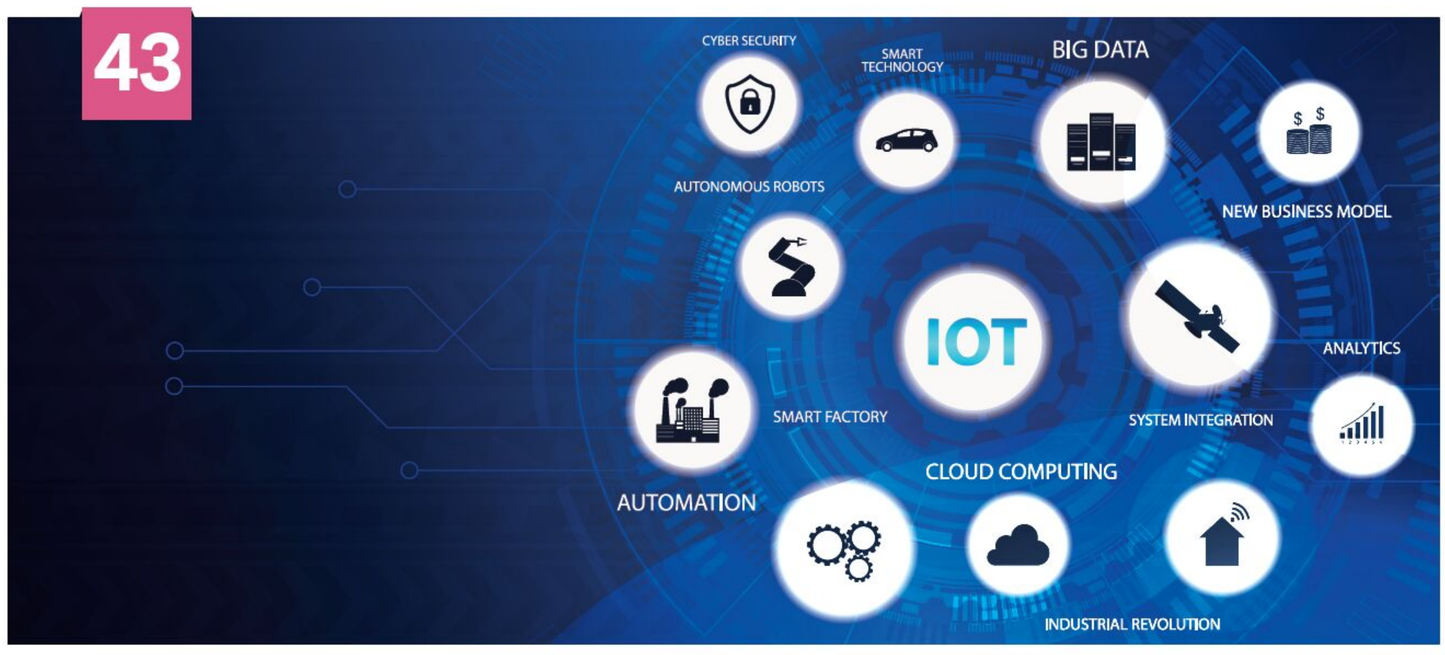
96 Couchbase: The Distributed Open Source Database



Open Source India 2023: Where The Industry Met The Community



Generative AI in IoT: Moving Towards Transformative Applications



The Role of Open Source in the Growth and Adoption of IoT

REGULAR FEATURES

07 FossBytes

EFYGROUP
YOURS SINCE 1969

Wanna Support Open Source For You?

Subscribe to the Magazine— so that we can keep promoting Open Source.



AMAZING OFFER

Pay for 6 Issues Get 12!

(Buy One Get One Free!)

WORLD'S LEADING PUBLICATION ON OPEN SOURCE

TEAR OUT — **ORDER FORM** — TEAR OUT

Please Tick (✓) Your Choice

☐ **Pay for 12 Issues Get 24 Issues**
and save 50%
(Buy 12 Issues Get 12 Issues Free!)
₹1440

☐ **Pay for 24 Issues Get 48 Issues**
and save 50%
(Buy 24 Issues Get 24 Issue Free!)
₹2880

☐ **Pay for 36 Issues Get 72 Issues**
and save 50%
(Buy 36 Issues Get 36 Issue Free!)
₹4320

To subscribe online, visit
https://tinyurl.com/y5kuv4la

OR SCAN THIS CODE

Name_____ Organisation_____ Mailing Address_____

_____ City_____

Pin Code_____ State_____ Phone No._____ Email_____

Subscription No. (for existing subscribers only)_____. **I would like to subscribe to the above (✓)marked Open Source For You magazine starting with the next issue.** Please find enclosed a sum of Rs_____ by DD/MO/crossed cheque bearing the No._____ dt._____ in favour of **EFY Enterprises Pvt Ltd,** payable at Delhi. *(Please add Rs 50 on non-metro cheque)*

Please mark one (nearest) relating to your subscription: ☐ **Indian Company** ☐ **MNC** ☐ **R&D organisation** ☐ **Engineering institute** ☐ **College/School** ☐ **Any other (specify):**_____

Send this filled-in form or its photocopy to : **EFY Enterprises Pvt Ltd,** D-87/1 Okhla Industrial Area, Phase 1, New Delhi 110 020 | Ph: **011-40596600** | e-mail: **support@efy.in**

Terms:- # These rates are applicable for new subscribers as well as renewal by existing subscribers. # Can access ezine till your subscription is active # The rates are valid for subscribers within India only. # Please allow 4-6 weeks for processing of your subscription. # The subscription copies will be dispatched through ordinary post only # Subscription Agents will not get agency commission against this scheme # Disputes, if any, are subject to exclusive jurisdiction of competent courts and forums in Delhi/New Delhi only. * Replacement will be made if intimation of damaged / non-receipt of copies is received within 30 days of its publication ** After three months, if you are not satisfied with the magazine, your balance amount will be returned (Not applicable for gift offer)

EDITOR
RAHUL CHOPRA

EDITORIAL, SUBSCRIPTIONS & ADVERTISING
Delhi (HQ)
D-87/1, Okhla Industrial Area, Phase I, New Delhi 110020
Phone: +91-9811155335
E-mail: info@efy.in

MISSING ISSUES
Phone: +91-9811155335
E-mail: support@efy.in

BACK ISSUES
Phone: +91-9811155335
E-mail: support@efy.in

NEWSSTAND DISTRIBUTION
Phone: +91-9811155335
E-mail: efycirc@efy.in

ADVERTISEMENTS
NEW DELHI (HEAD OFFICE)
Phone: +91-9811155335
E-mail: efyenq@efy.in

MUMBAI
E-mail: rmowest@efy.in

BENGALURU
E-mail: rmosouth@efy.in

CHINA
Worldwide Focus Media
E-mail: china@efy.in

GERMANY
pms Plantenberg Media Service GmbH
E-mail: germany@efy.in

JAPAN
Tandem Inc.
E-mail: japan@efy.in

TAIWAN
J.K. Media
E-mail: taiwan@efy.in

UNITED KINGDOM
ASA Media
E-mail: uk@efy.in

UNITED STATES
E & Tech Media
E-mail: usa@efy.in

Printed, published and owned by Ramesh Chopra. Printed at Tara Art Printers Pvt Ltd, A-46,47, Sec-5, Noida, on 28th of the previous month, and published from D-87/1, Okhla Industrial Area, Phase I, New Delhi 110020. Copyright © 2023. All articles in this issue, except for interviews, verbatim quotes, or unless otherwise explicitly mentioned, will be released under Creative Commons Attribution-NonCommercial 3.0 Unported License a month after the date of publication. Refer to http://creativecommons.org/licenses/by-nc/3.0/ for a copy of the licence. Although every effort is made to ensure accuracy, no responsibility whatsoever is taken for any loss due to publishing errors. Articles that cannot be used are returned to the authors if accompanied by a self-addressed and sufficiently stamped envelope. But no responsibility is taken for any loss or delay in returning the material. Disputes, if any, will be settled in a New Delhi court only.

SUBSCRIPTION RATES			
Year	**Newstand Price (₹)**	**You Pay (₹)**	**Overseas**
Five	7200	4320	—
Three	4320	3030	—
One	1440	1150	US$ 120

Kindly add ₹ 50/- for outside Delhi cheques.
Please send payments only in favour of EFY Enterprises Pvt Ltd.
Non-receipt of copies may be reported to **support@efy.in**—do mention your subscription number.

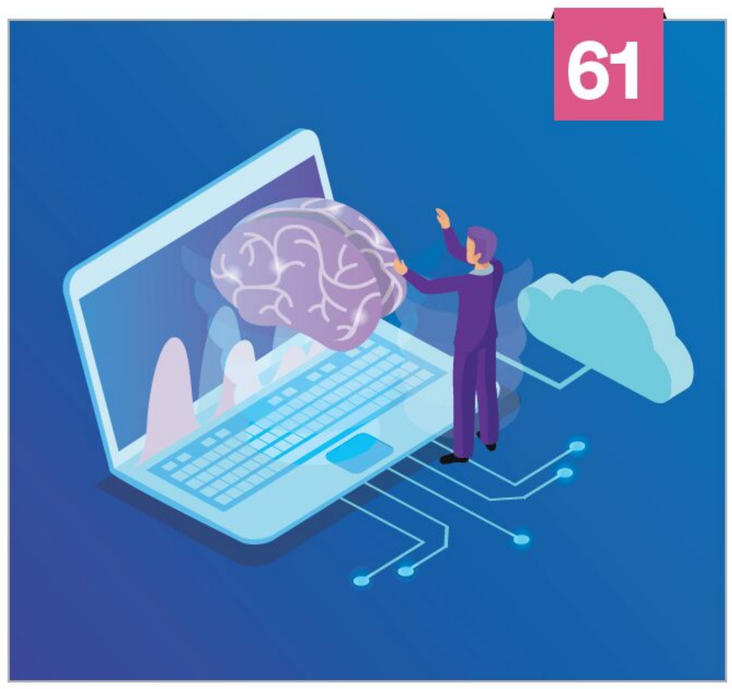
CONTENTS



How Open Source Tools Help with Real-time Data Analytics in IoT



The Future of IoT



Neuromorphic Computing: The Future of IoT



Simplifying the World of Distributed Ledger Technology



Latest Trends in DevOps and What's Next

Protect AI launches open source solutions for AI/ML security



Protect AI, known for its AI and machine learning (ML) security expertise, has introduced three open source software (OSS) tools—NB Defense, ModelScan, and Rebuff. The aim is to address the pressing need for security in the AI/ML field.

Amidst the proliferation of AI and ML tools, security has been a growing concern, and Protect AI aims to provide practical solutions. CEO Ian Swanson stated, "Our goal is to raise awareness for the need to make AI safer and provide tools organisations can start using immediately to protect their AI/ML applications."

NB Defense is a tool that scans Jupyter Notebooks—a key platform for data scientists. It detects leaked credentials, PII (personally identifiable information) disclosure, licensing issues, and security vulnerabilities within these documents.

ModelScan ensures the security of ML models shared online. It scans for unsafe code, safeguarding users who work with popular ML frameworks like PyTorch, TensorFlow, Keras, Sklearn, and XGBoost.

Protect AI's Rebuff addresses security concerns in large language models (LLMs). Vulnerable to prompt injection (PI) attacks, Rebuff offers a self-hardening prompt injection detection framework with four layers of defence.

These tools can be used as standalone solutions or integrated into the Protect AI platform, providing visibility, audibility, and enhanced security for ML systems.

Malicious npm package deploys rootkit

A malicious npm package that deploys an open source rootkit has raised concerns within the cybersecurity community. It marks a disconcerting trend, as it is the first known instance of a rogue package delivering rootkit functionality.

The malicious package, named 'node-hide-console-windows', cleverly mimics the legitimate npm package 'node-hide-console-window', exploiting minor typographical errors in what's known as typosquatting. Alarmingly, this fraudulent package amassed 704 downloads over two months before being removed from the npm registry.

The discovery was made by ReversingLabs, a software supply chain security firm, which first detected these activities in August 2023. According to ReversingLabs, the package facilitated the download of a Discord bot, which,

RISC-V warns export controls may stifle tech innovation

The chief executive of RISC-V International, Calista Redmond, has voiced concerns over potential US government restrictions on open source technology that could hamper chip development, impacting the global tech industry.

US lawmakers are urging the Biden administration to impose export control restrictions on RISC-V, an open source technology managed by RISC-V International. This technology is vital for creating chips in smartphones and AI.

RISC-V has found favour not only with major US firms like Qualcomm and Google but also with Chinese giant Huawei, raising national security concerns. Redmond defended RISC-V, comparing it to other open technology standards like Ethernet, which enables internet-connected computers to communicate.

Redmond warned that unprecedented restrictions on open standards might limit access to global markets, slowing product and talent flow. She cautioned against bifurcation at the standards level, which could lead to incompatible solutions and market closures.

Crucially, Redmond stressed that RISC-V has global contributions, from North America, Europe, and Asia. The standards published by RISC-V International do not offer a complete chip blueprint or proprietary information. The key difference lies in the freedom to use these open standards without requiring proprietary licences.

Redmond said, "Having access to open standards allows companies to innovate faster and spend their time creating differentiated products, rather than trying to reinvent the wheel."

Microsoft rushes security updates to address two vulnerabilities

Microsoft has issued urgent security updates for its products, including Edge, Teams, and Skype, to address two zero-day vulnerabilities linked to open source libraries. These vulnerabilities could result in crashes or, more concerning, arbitrary code execution.

One of the flaws, identified as CVE-2023-4863, stems from a heap buffer overflow issue in the WebP code library (libwebp). This vulnerability affects a range of browsers, including Safari, Mozilla Firefox, Microsoft Edge, Opera, and native Android web browsers, as well as apps like 1Password and Signal.

The second vulnerability, known as CVE-2023-5217, arises from a heap buffer overflow weakness in the VP8 encoding of the libvpx video codec library. This library plays a crucial role in VP8 and VP9 video encoding and decoding for desktop video players and streaming services like Netflix, YouTube, and Amazon Prime Video.

Microsoft acted swiftly to patch these vulnerabilities, focusing on products like Microsoft Edge, Microsoft Teams for Desktop, Skype for Desktop, and WebP Image Extensions. For users of WebP Image Extensions from the Microsoft Store, the security update will be automatically applied unless automatic updates are disabled.

These vulnerabilities were exploited in undisclosed spyware attacks, as reported by Apple Security Engineering and Architecture (SEAR), Google Threat Analysis Group (TAG), and the Citizen Lab. While details of these attacks remain elusive, it's clear that these vulnerabilities posed a real-world threat.

As cyber threats evolve, prompt software updates are essential. Microsoft's swift response demonstrates the commitment of tech giants to safeguard their products and users from emerging risks.

in turn, initiated the deployment of an open source rootkit known as 'r77'. This discovery raises concerns about the potential use of open source projects as a conduit for malware distribution.



The malevolent code concealed within 'node-hide-console-windows' resides in its *index.js* file. Upon execution, it fetches and automatically runs an executable known as 'DiscordRAT 2.0', a C#-based open source Trojan with over 40 commands that enable remote attackers to take control of a victim's system via Discord while disabling vital security software.

One troubling command is '!rootkit', which activates the r77 rootkit on the compromised system. The r77 rootkit, maintained by bytecode77, is a 'fileless ring 3 rootkit' designed to conceal files and processes, either as a standalone tool or bundled with other software.

What sets this package apart is its reliance on publicly available components, providing low-stakes hacking groups with a straightforward path to launching supply chain attacks. This incident underscores the importance of rigorous vetting and verification in the open source ecosystem.

Critical 'ShellTorch' flaws found in TorchServe

Researchers have identified a set of critical vulnerabilities within TorchServe, a machine learning framework, which could potentially allow cyberattackers to compromise artificial intelligence (AI) models, leading to various undesirable outcomes. This revelation underscores the susceptibility of AI applications to open source bugs, even affecting tech giants like Google and Amazon.



TorchServe, maintained by Amazon and Meta, is an open source framework that deploys deep learning models using the PyTorch library. It is widely utilised by major companies such as Walmart, Amazon, Microsoft Azure, and Google Cloud.

Exploitation of these vulnerabilities could grant threat actors access to proprietary AI model data, enabling them to insert malicious models into production environments, manipulate machine learning results, and gain control over servers.

What's concerning is the widespread exposure of these vulnerabilities. Thousands of instances of vulnerable TorchServe software are accessible on the internet, posing a threat of unauthorised access and malicious actions, as reported by Oligo, the researchers behind the discovery.

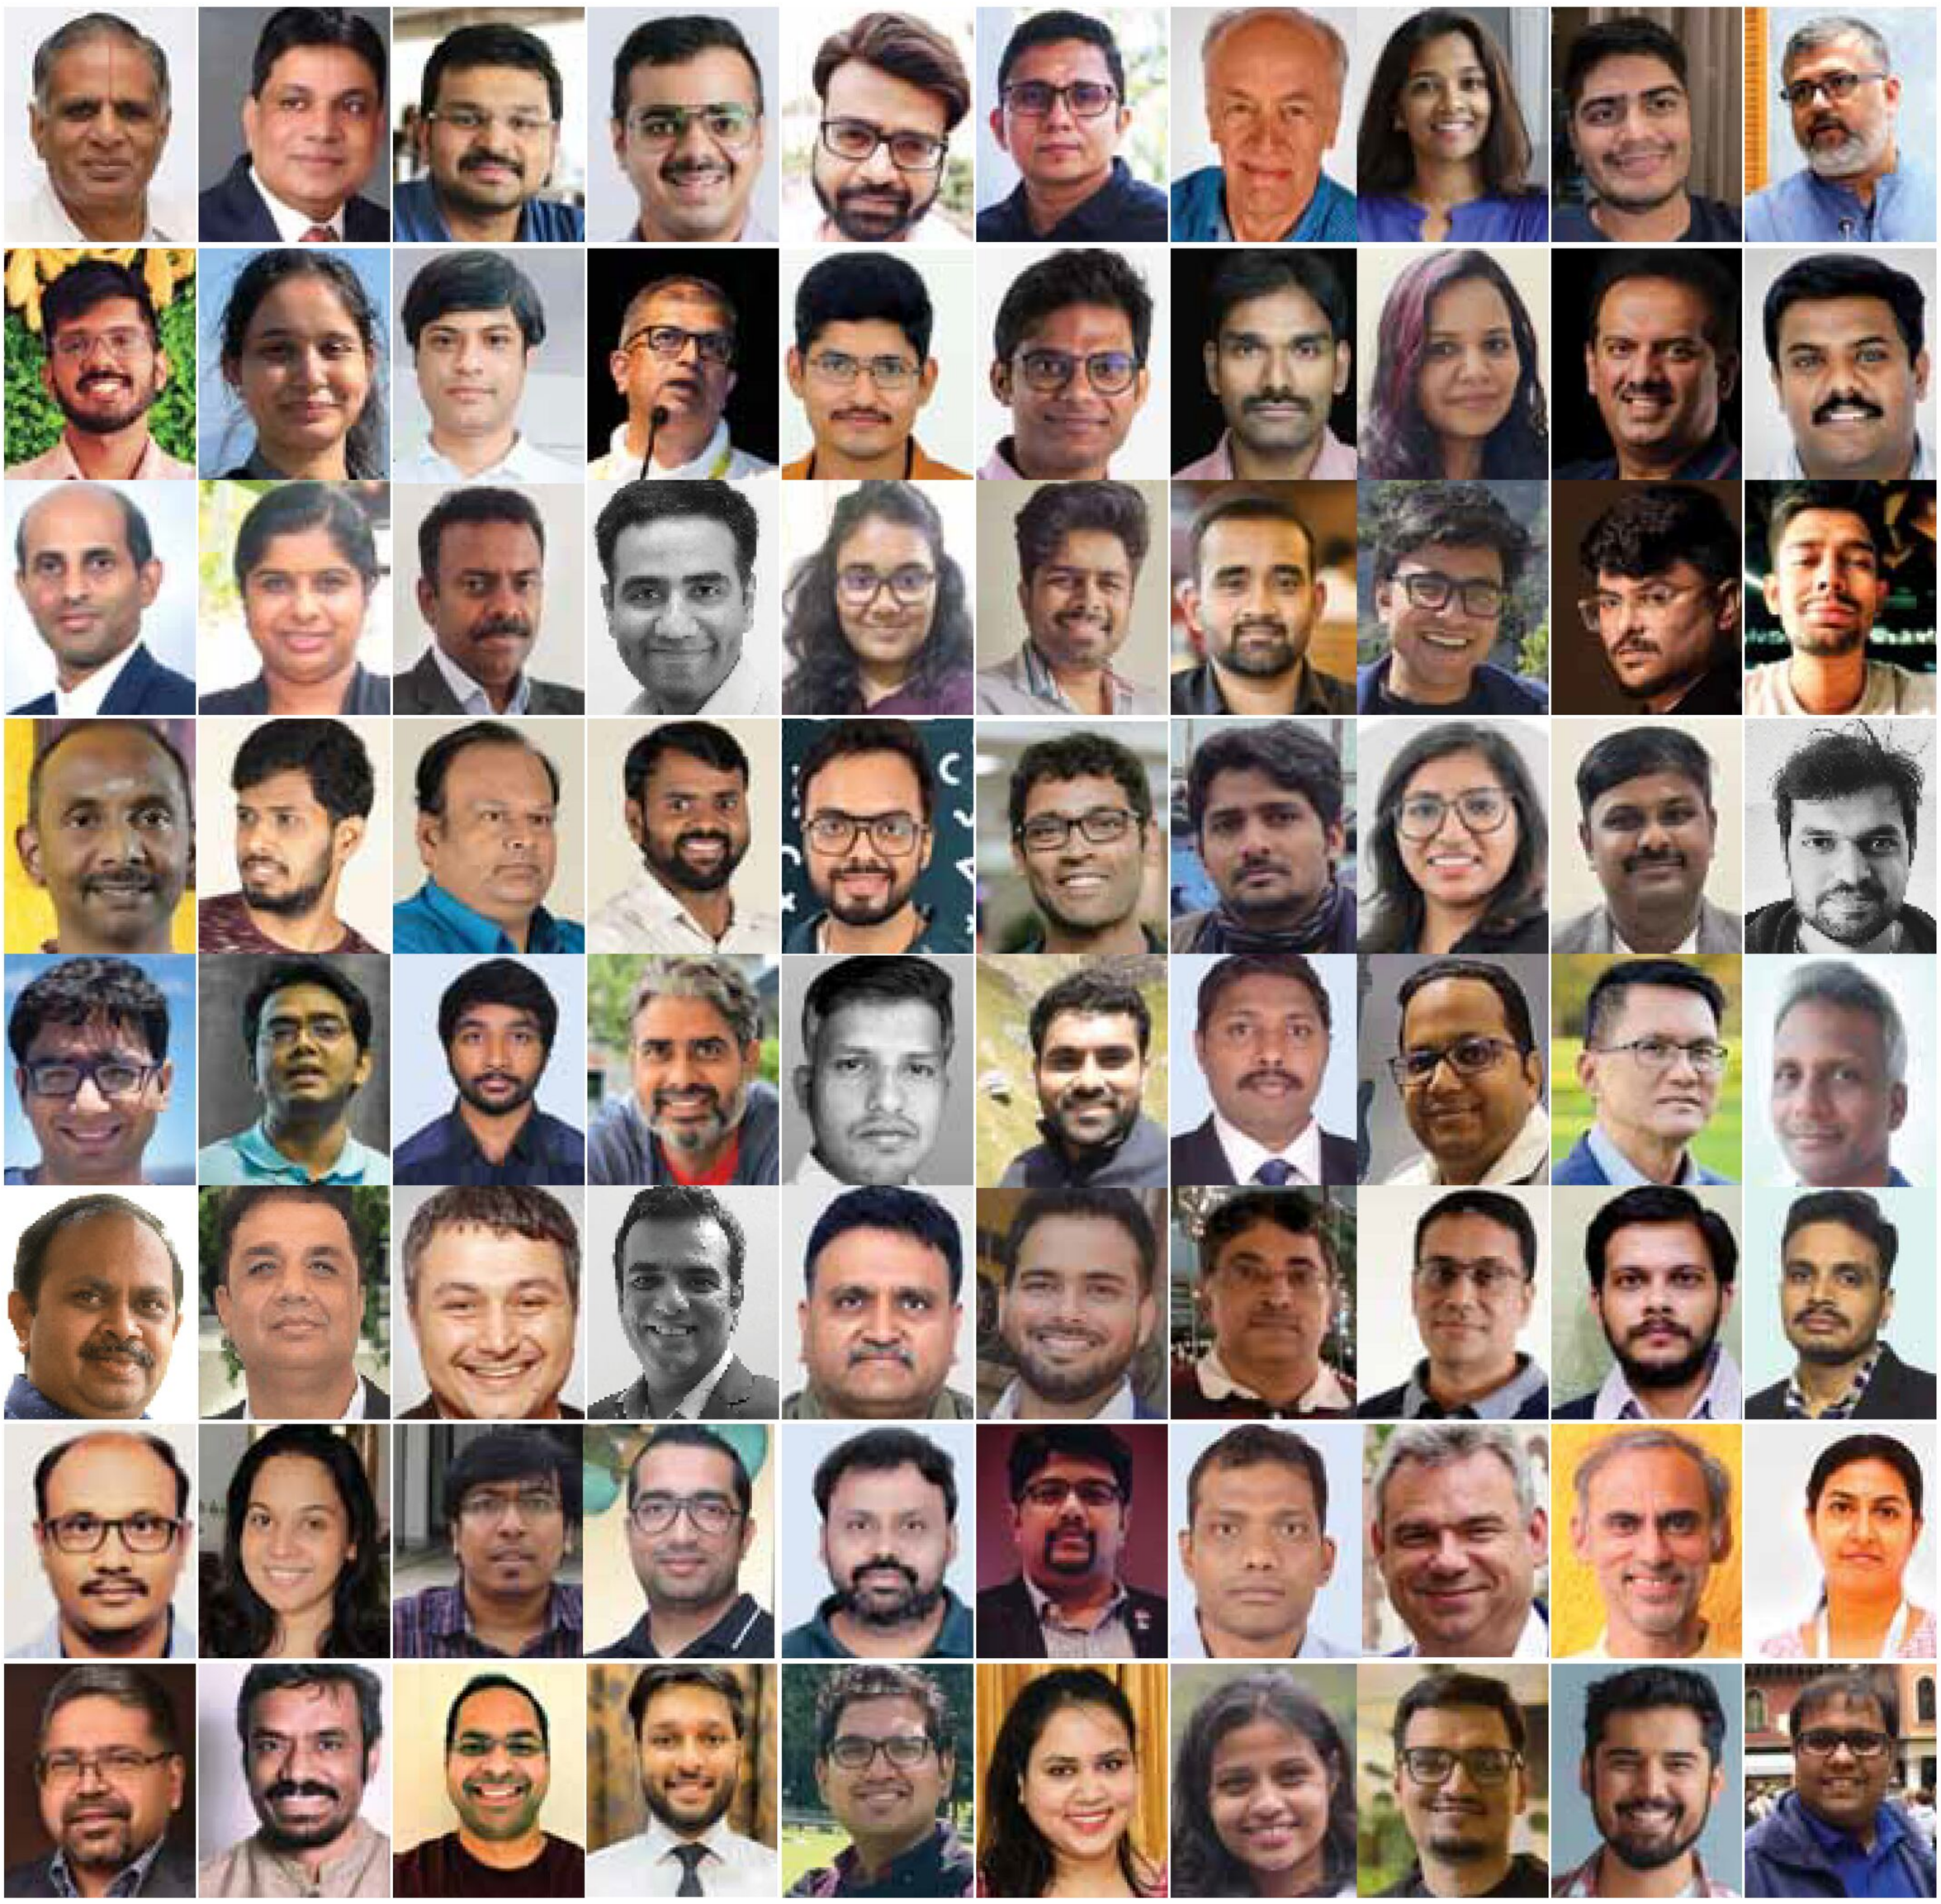
www.OpenSourceIndia.in

20th Edition

OPEN SOURCE INDIA

THANK YOU

Speakers



For more details, call on +91-98111-55335 or email us at info@opensourceindia.in

Furbinator 3000: This DIY camera could be a nature photographer's new ally

IT expert James Milward has harnessed artificial intelligence to transform an outdoor camera, which he calls Furbinator 3000, into a sentinel for his garden, proficient at identifying and deterring unwanted badger and fox visitors. However, this novel project, initiated as a solution to keep these creatures at bay, has caught the attention of nature photographers, who perceive its potential to revolutionise their craft.

Milward's ingenious endeavour entails custom-built code, closely monitoring a Ring camera's live feed to discern whether the guests are the bothersome creatures he seeks to deter. When the camera spots a badger or fox, it activates an ultrasonic repellent he procured from Amazon. The intriguing prospect here is how this technology can be adapted for photographers to enhance their wildlife photography.

The core of Milward's project centres on training the cameras to distinguish between foxes and badgers, as each requires a different frequency of sound for effective deterrence. By enabling the camera to identify the intruder, the right deterrent can be activated, offering a promising addition to wildlife photographers' toolkits.

To build this innovation, Milward employed a Raspberry Pi 4 and TensorFlow Lite, a renowned machine learning tool. Video clips from his Ring camera were manually tagged as foxes or badgers using the LabelImg tool. This critical step prevented the computer from misidentifying common objects. The model was refined with additional steps by utilising Google Colab.

All in all, this project offers a practical solution for photographers seeking to perfect their craft while coexisting with garden wildlife.

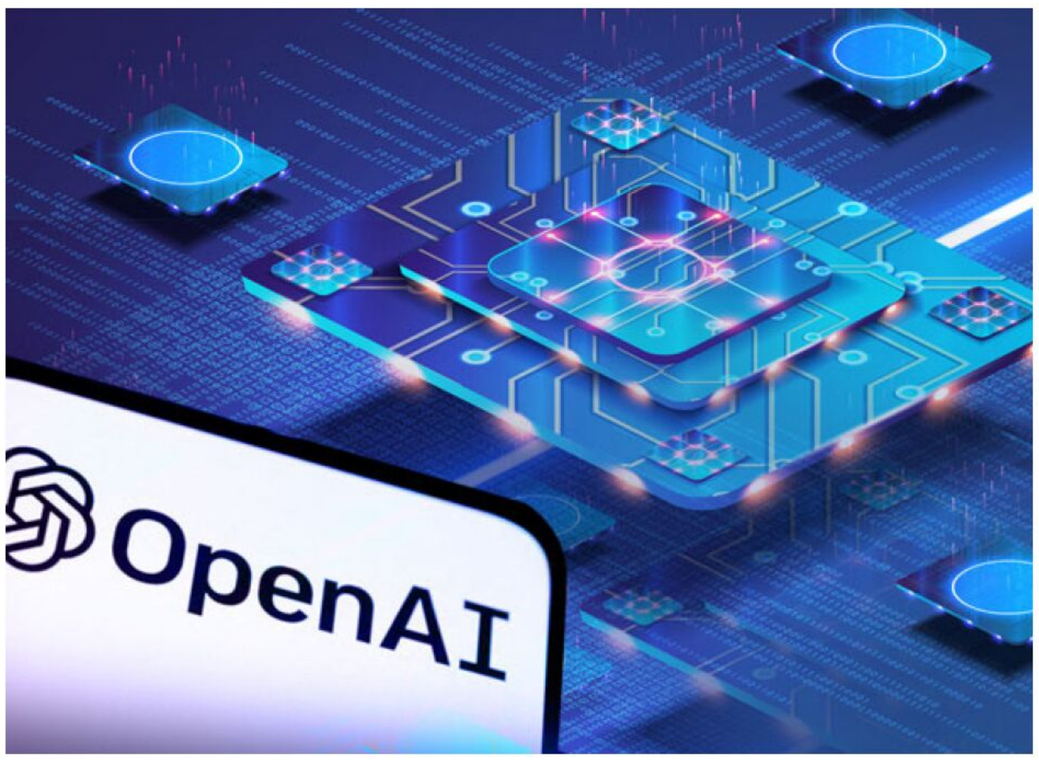
Oligo has collectively termed these vulnerabilities as 'ShellTorch'. Two of these vulnerabilities are rated critical, with near-maximum severity ratings on the CVSS scale. They include CVE-2023-43654, a server-side request forgery (SSRF) vulnerability allowing remote code execution, and CVE-2022-1471, a Java deserialisation RCE flaw.

The third vulnerability arises from TorchServe's default configuration, which exposes a critical management API to the internet. Although modifying this configuration can mitigate the issue, many organisations and projects using TorchServe have retained the default settings.

These vulnerabilities highlight the reality that AI applications face the same risks as other software when it comes to open source code. However, in AI, the consequences can be severe due to its numerous applications.

OpenAI seeks hardware independence

Nvidia currently dominates the AI market, providing the hardware essential for training expansive artificial intelligence (AI) language models. With the rising demand, the availability and prices of their chips have become challenging.



OpenAI, the creator of ChatGPT, is actively exploring alternatives to Nvidia's solutions, including the possibility of developing its proprietary chips.

The company is devising plans to address the current hardware limitations in the AI sector. Among the solutions under consideration, one radical approach includes acquiring one of the former rivals and granting it the capability to manufacture its chips. So far, the investigation has reached the stage of identifying potential merger candidates.

Sam Altman, OpenAI's CEO, has underscored the urgency of procuring more AI chips, marking it a paramount concern for the company. Yet, fulfilling this need isn't straightforward. Nvidia is grappling with manufacturing constraints, resulting in an inability to meet the demand for its H100 AI chips for an anticipated 1.5 years. The ChatGPT creator previously harnessed 10,000 GPUs to train ChatGPT. Its ambition to scale further is hampered by the scarce supply and skyrocketing prices of chips.

The financial implications are staggering. For the company to handle a query volume equivalent to a tenth of Google's, a whopping $48 billion investment in graphics processing units (GPUs) would be necessitated. Moreover, maintaining this scaled operation would subsequently demand an annual expenditure of $16 billion.

The company's contemplation of custom chips mirrors strategies adopted by tech behemoths like Amazon, Google, and Meta, who have bespoke chips tailored to their unique requirements. Microsoft, the OpenAI's principal partner, is reportedly developing its custom silicon. However, its journey to custom silicon realisation will span years. Meanwhile, the reliance on chip giants like Nvidia and AMD remains inevitable.

OpenUK and UN unite for Open Source Data Centre Challenge

In a bid to combat the environmental impact of data centres, OpenUK, representing the UK's open technology sector, has partnered with the United Nations' Office

of Information and Communications Technology (UN-OICT) to launch the 'Open Source Data Centre Challenge'. This project aims to develop open source solutions for creating a carbon-negative data centre model.



Data centres, responsible for around 3-4% of global carbon emissions, have a larger carbon footprint than aviation, as per 2022 data. The initiative, known as 'Patchwork Kilt', will bring together a diverse community to tackle this issue. Participants in the challenge will contribute to areas such as heat redirection, building repurposing, heat reuse, circular supply chain management, blueprint evolution, and community empowerment.

Christ Lloyd-Jones, OpenUK's chief blueprint officer, said, "Our collaboration creates a challenge to open this initiative to contributions and a broader community, allowing new technologies to be included. It's a critical initiative that can impact data centre sustainability."

Salem Avan, UN director of policy, strategy, and governance, highlighted open source's importance, stating, "It fuels innovation, transparency, and cost-efficiency. By using open source technology to run data centres efficiently, we reduce our environmental footprint and contribute to global sustainability."

The winning contributor will receive a trophy and attend OpenUK's Awards Ceremony in November.

Vulnerabilities in Curl raise concerns

Maintainers of a widely-used open source tool, essential for supporting network protocols such as SSL, TLS, HTTP, FTP, and SMTP, have cautioned users about two upcoming vulnerabilities. Curl, a command-line utility for API interaction and file downloads, faces one high-severity issue (CVE-2023-38545) and a lower-severity problem (CVE-2023-38546).

While CVE-2023-38545 affects both Curl and its underlying library, libcurl, CVE-2023-38546 exclusively affects libcurl. An update released on October 11 addressed both issues.

Maintainers have been cautious about disclosing specific version details, emphasising the need to avoid premature exploitation.

Melissa Bischoping, director of endpoint security research at Tanium, highlighted the widespread use of Curl. She advised organisations to take heed and assess their systems in light of these vulnerabilities, even though the full impact remains uncertain.

Warehouse: A new tool for easier Flatpak management

Flatpak enthusiasts can now rejoice as a user-friendly tool named 'Warehouse' will simplify the management of Flatpak apps and remotes. This GTK4/libadwaita-powered tool bridges the gap between Flatpak's command-line interface and user convenience.

Warehouse offers a quick overview of installed Flatpak apps and runtimes, with an option to reveal hidden runtimes. It matches the detail of the 'flatpak list' command but in a user-friendly GUI, providing buttons for convenient clipboard copying.

VLC unveils version 3.0.19

VLC, the renowned open source media player, has unveiled version 3.0.19, delivering substantial upgrades along with critical bug fixes. However, there's a caveat — the most significant improvements are exclusively reserved for Windows users.

VLC 3.0.19 introduces super resolution scaling for compatible Nvidia and Intel GPUs. This AI-assisted upscaling is best suited for lower-resolution videos, albeit not a game-changer for high-definition content. Windows users can enable this feature via VLC's *Advanced Settings > Output Modules* section. The update also bestows AV1 hardware decoding on Windows, a notable enhancement for video playback.

The positive news is that all platforms, including Linux, benefit from improved software decoding for AV1 HDR content.

The update also addresses several issues and enhances various functionalities. These include improvements in FFmpeg-muxed MP4 chapter control, AAC handling, playback of QNap-produced AVI files, and the surprising comeback of support for very old RealVideo format files. Additional adjustments cover rendering issues on Linux when using the VLC full screen controller and enhanced subtitle clarity for macOS users.

VLC 3.0.19 claims to rectify YouTube playback issues, provide better file handling during read-only mode, and offer improved CrunchyRoll SSA rendering.

Although VLC 3.0.19 has not yet been officially announced on the Videolan website, the source code is accessible, and installers for Windows and macOS are available. Linux users will need to wait for their turn to access these new features.

Raspberry Pi's Bookworm release adopts Wayland and PipeWire

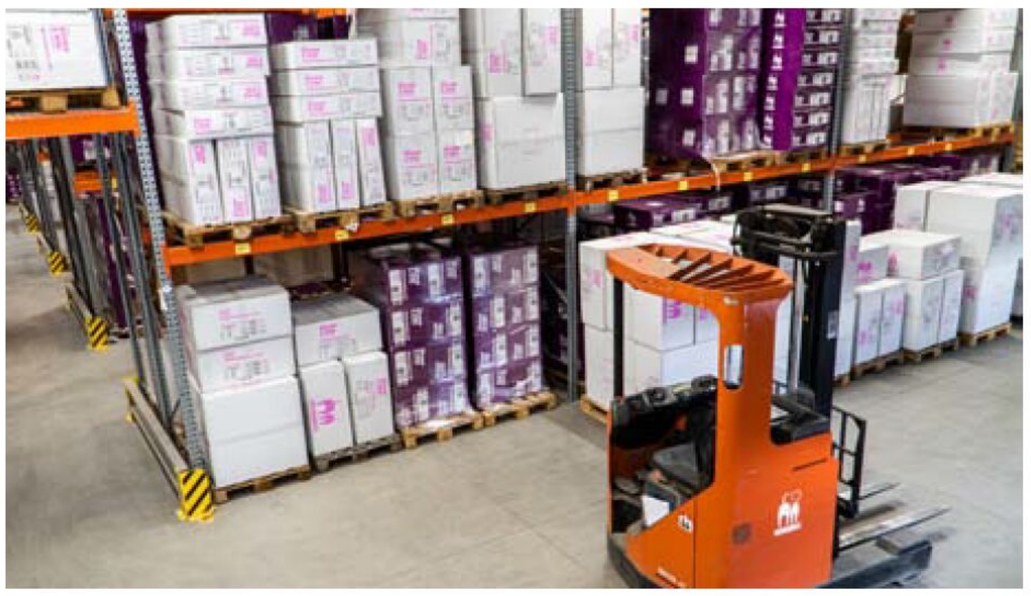
Raspberry Pi's latest OS release, codenamed 'Bookworm', offers a revamped experience, building upon the recent Debian 12 'Bookworm' release. Simon Long from Raspberry Pi explained, "We have been working on some major architectural changes to the Raspberry Pi Desktop, and these are launched for the first time in the Bookworm release."

The most significant change is the shift from X11 to Wayland as the display system, coupled with Wayfire instead of Mutter as a compositor. This transition aims to enhance OS performance, security, and add subtle window animations. The taskbar at the top is updated, now Wayland-native, with plugins from the previous setup and a new GPU graph plugin.

PipeWire takes over as the default sound system, replacing PulseAudio. NetworkManager replaces dhcpcd for networking, giving users more connectivity options. Despite the familiar desktop appearance, the underlying technology has evolved significantly.

A Raspberry Pi-optimised Firefox version has been introduced, with V4L2 codec support, Widevine DRM, and compatibility with CSI cameras via libcamera, alongside desktop sharing on Wayland.

In summary, Raspberry Pi OS Bookworm is packed with valuable improvements for enthusiasts. To get your hands on it, you cannot upgrade from the previous version, Raspberry Pi Bullseye. Instead, you must flash a Bookworm image to an SD card or SSD using the Raspberry Pi Imager app. Alternatively, download an IMG from the official Raspberry Pi website's *Downloads* page. For those who encounter issues with Wayland, Wayfire, and PipeWire, the *Advanced Settings* menu in raspi-config offers options to revert to X11/Openbox and PulseAudio.



Warehouse lets you clear user and app data for installed Flatpaks, offering a seamless refresh option without uninstalling. It also scans for residual data left by uninstalled Flatpak apps. Users can selectively remove this data per app or opt for a one-click solution, effectively reclaiming storage space.

Managing Flatpak remotes is simple with Warehouse, as new remotes can be easily removed or added through a graphical interface, simplifying the process.

In essence, Warehouse emerges as a valuable tool for Flatpak enthusiasts, offering streamlined management without the need for command-line operations.

Ubuntu halts release 23.10 due to hate speech in Ukrainian translations

Ubuntu has taken the unprecedented step of suspending its Desktop release 23.10. This move comes after the discovery of hate speech within Ukrainian translations, which included anti-Semitic, homophobic, and xenophobic slurs. The offensive content was introduced through a 'third-party tool' external to the Ubuntu Archive.

The Ubuntu project issued a statement confirming the hate speech, attributing it to a malicious contributor who used a 'third-party online service' relied upon by the Ubuntu Desktop Installer for language support. The 23.10 image was promptly removed, with assurances that a corrected version will be available once the proper translations are reinstated.

Notably, this issue affects translations presented during installation via the Live CD environment. If you've upgraded to Ubuntu Desktop 23.10 from a previous release, you are unaffected.

The impacted images were Ubuntu Desktop 23.10 and Ubuntu Budgie 23.10, while the Ubuntu desktop legacy ISO remains unaffected.

Translations are fundamental to internationalising applications, supported by contributions from individuals globally. Canonical and Ubuntu strongly disapprove of hate speech and offensive language, in line with their code of conduct.



A GitHub pull request removed the 'insulting [localisation] strings' around October 12. BleepingComputer observed that these malicious Ukrainian strings were cleverly concealed at the end of the translations file.

Ubuntu has restored Ukrainian translations and is conducting a comprehensive audit before official release. Meanwhile, users are advised to download Ubuntu Desktop 23.10 from the Ubuntu *Downloads* page using the unaffected desktop legacy installer ISO or upgrade from a previously supported Ubuntu release.

OpenSSF unveils repository to combat malicious open source packages

Open Source Security Foundation (OpenSSF) is taking on malicious open source software by launching the 'Malicious Packages' repository. The aim is to streamline open source security and counter potential threats.

Caleb Brown, a senior software engineer at the Google Open Source Security Team, and Jossef Harush Kadouri, head of software supply chain security at Checkmarx, highlighted the current complexities in dealing with malicious open source packages. Multiple open source package repositories handle such packages in unique ways, often leading to a convoluted process. Typically, when the community reports a malicious package, the package repository's security team removes it along with its metadata. Unfortunately, these actions usually occur without a trace or public record.



This fragmented approach has made identifying and tracking malicious packages a challenge, necessitating extensive data gathering from various public sources or reliance on proprietary threat intelligence feeds.

OpenSSF's Malicious Packages repository acts as a central hub, where reports of malicious open source packages are stored and accessible to the wider community. The repository simplifies the process of tracking, identifying, and countering threats within the open source ecosystem.

Brown and Kadouri emphasised the repository's significance, stating, "OpenSSF believes that having a public repository of this information will 'stop malicious dependencies from moving through CI/CD pipelines, refine detection engines, scan for and prevent usage in environments, or accelerate incident response'."

A noteworthy feature of the repository is its use of the open source vulnerability (OSV) format for report storage, ensuring compatibility with various tools such as the osv.dev API, the osv-scanner tool, and deps.dev.

The repository compiles data from diverse sources, including Checkmarx security, exports of malicious packages tracked by GitHub, and the Package Analysis project. This detailed behavioural analysis helps to determine whether a package displays malicious traits.

AMD to acquire Indian-founded Nod.ai, bolstering AI portfolio

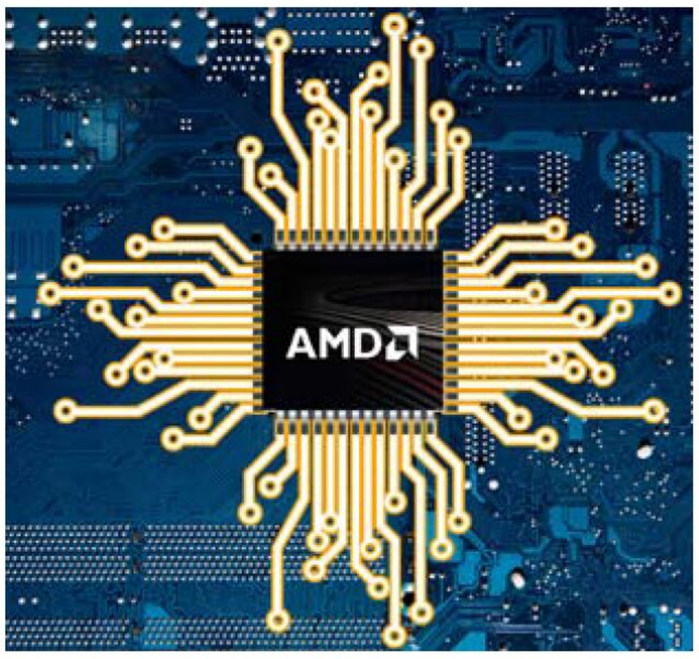
Advanced Micro Devices, Inc. (AMD), the Californian semiconductor giant, has sealed the deal to acquire Nod.ai, an open source AI software startup co-founded by Anush Elangovan and Harsh Menon, marking a pivotal step towards enhancing its foothold in the AI domain.

Nod.ai, established in 2013, has been making waves with its expertise in machine intelligence, simplifying the training and deployment of machine learning models using advanced compiler and code generation technologies.

Vamsi Boppana, senior vice president of AMD's Artificial Intelligence Group, emphasised the significance of this acquisition, stating, "The acquisition of Nod.ai is expected to significantly enhance our ability to provide AI customers with open software that allows them to easily deploy highly performant AI models tuned for AMD hardware."

Nod.ai's software, notably SHARK, is designed to automate the deployment of high-performance AI models across various platforms, effectively reducing the need for manual optimisation. Anush Elangovan, CEO of Nod.ai, underscored their commitment to addressing the ever-evolving challenges in the industry. "By joining forces with AMD, we will bring this expertise to a broader range of customers on a global scale," he said.

This strategic move aligns with AMD's commitment to building a comprehensive open software ecosystem, facilitating the deployment of AI solutions. Nod.ai's technology is already widely utilised in cloud, edge, and endpoint devices.



The acquisition of Nod.ai signifies a significant development in AMD's AI growth strategy, underlining their focus on making AI solutions more accessible through developer tools, libraries, and models.

This landmark deal comes a year after AMD's acquisition of Xilinx, Inc., for $50 billion, solidifying its position in the semiconductor industry.

For more news, visit *www.opensourceforu.com*





Asia's #1 Open Source Conference

Total registrations

15,372

Total attendees

4743

Speakers

80









www.OpenSourceIndia.in

An EFY GROUP EVENT

20th Edition

OPEN SOURCE INDIA

Thank You

VISITORS, for making the event a huge success

SPEAKERS, for contributing your valuable thoughts

PARTNERS, for your support

intel | OpenVINO

aws

Clari

new relic.

salesforce

ManageEngine Site24x7

soda foundation

PURESTORAGE

GitHub

thoughtworks

BOEING

Mydbops
Scaling Database Operations

SIEMENS

rezoomex

Findy

DELL

CODE AND THEORY

mt-b

JAVASCRIPT MEETUP
(BANGALORE, INDIA)

Helical Insight

soda foundation
meetupindia

MOSIP

CAZE LABS
ignite innovations, harness open source!

Waterdip

ptimus

skydo

For more details, call on 98111 55335 or email us at info@opensourceindia.in

PDF Tricks for the Linux Command-Line

There are a myriad of PDF utilities, each with its own idiosyncrasies. This article combines several of them and provides just the right syntax to get you going without any delay.



In the January 2023 edition of *OSFY*, I wrote an article on how to programmatically create HTML, ODT, DOCX and PDF files. In this article, I focus on PDF so that you can use it as a barebones ready-reference for your oft-required PDF-processing tasks.

Creating PDFs from scratch

Old MarkDown is passé. CommonMark is the new standard, or rather, it is standardised MarkDown. If you prefer MarkDown, then you have to use its slow Perl script. If you can use CommonMark, then get the executable from my website. I built it from its C source code, and it is blisteringly fast.

```
# Convert MarkDown to HTML
perl markdown.pl jokebook.md > content.htm

# or
cmark --unsafe --validate-utf8 \
      jokebook.md > content.htm

# Place the converted HTML in a HTML template
echo '<!DOCTYPE html><html><title>2020 Jokebook</title></
head><body>' > jokebook.htm
cat content.htm >> jokebook.htm
echo '</body></html>' >> jokebook.htm

# Embed images in the HTML
libreoffice \
    --convert-to "html:HTML:EmbedImages" \
    jokebook.htm
# creates jokebook.html with self-contained
# (base64-encoded) images

# Convert HTML to ODF
libreoffice --convert-to "odt" jokebook.html
# creates jokebook.odt

# Convert ODF to PDF
libreoffice --convert-to "pdf" jokebook.odt
# creates jokebook.pdf
```

Converting images to PDF

Sometimes, you have to create PDF pages from images. ImageMagick is the preferred Linux utility to convert images.

```
# Convert images to PDFs
magick front-cover.png -resize 100% front.pdf
magick back-cover.png -resize 100% back.pdf
```

Concatenating several PDFs

PDFtk is a powerful PDF-processing utility that can perform tasks such as merging, splitting, encrypting, decrypting, stamping, and watermarking.

```
# Concatenate several PDFs
pdftk front.pdf inner-pages.pdf back.pdf \
      output book.pdf
```

Encrypting PDFs

PDF documents can be encrypted using two passwords — *owner* and *user.*

```
pdftk book.pdf output book-encrypted.pdf \
      encrypt_128bit \
      owner_pw RcHrDsTlMn^012 \
      user_pw FrSfTWrFnDtn^321
```

If you set an empty user password, you can let users view the PDF without a password prompt. You can specify additional restrictions using the 'allow' option. Possible values are *Printing, DegradedPrinting, ModifyContents, Assembly, CopyContents, ScreenReaders, ModifyAnnotations, FillIn* and *AllFeatures.* If you do not specify the *allow* option, none of these features will be available. Well, the PDF standard specifies that these features should not be available. In reality, the restrictions are not strictly implemented by many PDF viewer applications.

Removing PDF password

Use the *input_pw* option to specify the password.

```
pdftk book-encrypted.pdf \
      input_pw RcHrDsTlMn^012 \
      output book-decrypted.pdf
```

If you start this command with a space, then the password (RcHrDsTlMn^012) will not be stored in the 'bash' shell history. Another trick is to use an interactive dialog.

```
sPassword=$(zenity --password \
           --title "Decrypt PDF" \
           --text "Type the password")
pdftk book-encrypted.pdf \
      input_pw $sPassword \
      output book-decrypted.pdf
```

If you want a console-only prompt, then disable the input echo.

```
stty -echo
read -p "Type the password: " sPassword
stty echo
```

Converting PDF pages to images

You can use the *pdftoppm* utility to convert PDF pages to images in JPEG and PNG format. You can also specify the pixel density and page range.

```
# Export pages 2 to 12 with 96 dpi
pdftoppm -png -r 96 -f 2 -l 12 book.pdf page
# creates numbered images with prefix 'page'
```

Rasterising a PDF

Sometimes, you cannot give a PDF as is to someone else or put it online. To foil content scrapers (and since 'AI' is now an unabashed content scraper), it is best to convert even the text to JPEG images. Why JPEG? Because it is lossy.

```
pdftoppm -jpeg -r 96 book.pdf page
magick page*.jpg book-rasterized.pdf
```

Some content scrapers do have the ability to read text from images, but that task requires extra resources. If you set the DPI to 72, it gets more demanding. For further annoyance, stamp it or watermark it.

```
magick watermark.png watermark.pdf
pdftk book-rasterized.pdf \
      stamp watermark.pdf \
      output pages-watermarked.pdf
```

Converting PDF to DjVu

Many PDF books on *Archive.org* are scanned from library copies of physical books. These PDFs are extremely heavy as the scans are usually dense or big images. Ordinary PDF readers on tablets struggle to read such PDFs. My solution is to convert the PDFs to DjVu. (DjVu viewer applications are optimised for reading from images. They load DjVu pages in a snap.)

```
pdf2djvu --dpi=220 \
         --output=tablet.djvu \
         library.pdf
```

Depending on the page size in the PDF, you will have to adjust DPI with the --dpi option.

Continued to page....19

Fintech: The Digital Evolution of Finance

Fintech is reshaping finance, ushering in a new era of digital prosperity, and changing the way we interact with money.

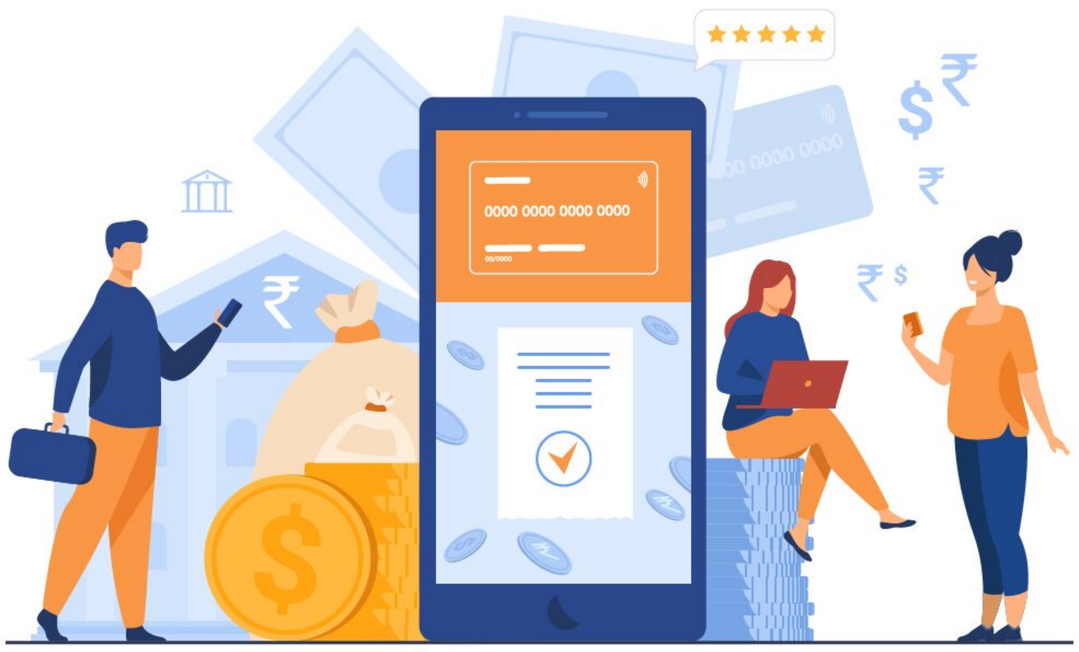


Image Source: https://www.freepik.com

Fintech encompasses leveraging and integrating the financial services technology provided by various organisations to offer better customer services. This marks a fundamental shift in the way banking and financial services were previously delivered. It introduces innovative services, embodying a digital transformation philosophy through the adoption of easy and interoperable solutions and services. Initially conceived as backend technology, fintech has evolved to become more customer-centric and now forms the front-end of customer operations. It aims to provide very low-cost transactions and near real-time intelligence in financial transactions.

Key business solutions from fintech

Fintech has brought disruptive changes to the banking and financial services industry, impacting various key areas.

- ***Digital payments and mobile banking:*** Payments made using mobile apps eliminate the need for explicit interactions with banking systems
- ***Voice payments:*** Speech-to-text interfaces are being used for making payments
- ***Virtual cards:*** Payments are being made using virtual credit and debit cards for e-commerce and online transactions, replacing physical cards
- ***Decentralised finance (DeFi):*** This is being used to eliminate/minimise service fee and charges
- ***Crowdfunding:*** Small investments are being solicited from a broad group of investors
- ***Lending and other bank-related transactions:*** Innovative lending solutions are on offer

Key technologies used by fintech companies

Fintech relies on various financial technologies.

- ***Blockchain:*** Creating distributed ledgers for enhanced transparency
- ***Robotic process automation:*** Automating processes to handle repetitive tasks
- ***Big Data:*** Collecting and processing structured, semi-structured and unstructured data in real time, providing valuable insights
- ***Artificial intelligence/machine learning:*** Leveraging artificial intelligence and machine learning algorithms for better business results
- ***Cloud enablement with software-as-a-service and serverless technologies:*** Enabling rapid scalability by leveraging cloud services
- ***Quantum computing:*** Harnessing the high speed of quantum computers for cybersecurity, fraud detection, risk management, etc
- ***Open APIs and open source technologies:*** Connecting third-party APIs to financial systems
- ***Rapid application development:*** Leveraging low-code and no-code tools and techniques

Open source libraries for fintech

Fintech benefits from various open source projects, with some of the popular ones being:

- ***Qlib***
  - An open source project from Microsoft
  - A platform for quantitative investment solutions that leverages AI/ML models and market dynamics paradigms
- ***FinGPT***
  - A financial LLM (large language model) from AI4 Finance Foundation

Key benefits of financial technologies

Fintech offers a number of benefits, including:

- Exponential increase of transaction speed
- Real-time predictive models to act with very low latency and prevent/reduce failures
- Improved customer-centricity
- Higher user engagement with innovative solutions and services

Common disadvantages of fintech

Fintech also presents some disadvantages, such as:

- Privacy and security risks for users.
- Increased risks of data theft and identity frauds
- Cannibalisation and intense global competition leading to rapid disruptions
- Legal and regulatory challenges due to incoherent market activities

Fintech promotes financial inclusion by offering new opportunities to diverse communities that were previously excluded from traditional banking services. It leverages cutting-edge solutions to bring the unbanked into the banking industry. It fosters innovation in the financial sector, and extends financial benefits to the most underprivileged in the society, contributing to the greater global good. END

References

- *https://github.com/microsoft/qlib*
- *https://github.com/AI4Finance-Foundation*

By: Phani Kiran

The author has nearly 20 years of experience in the IT industry and has worked on digital transformation initiatives.

Continued from page....17

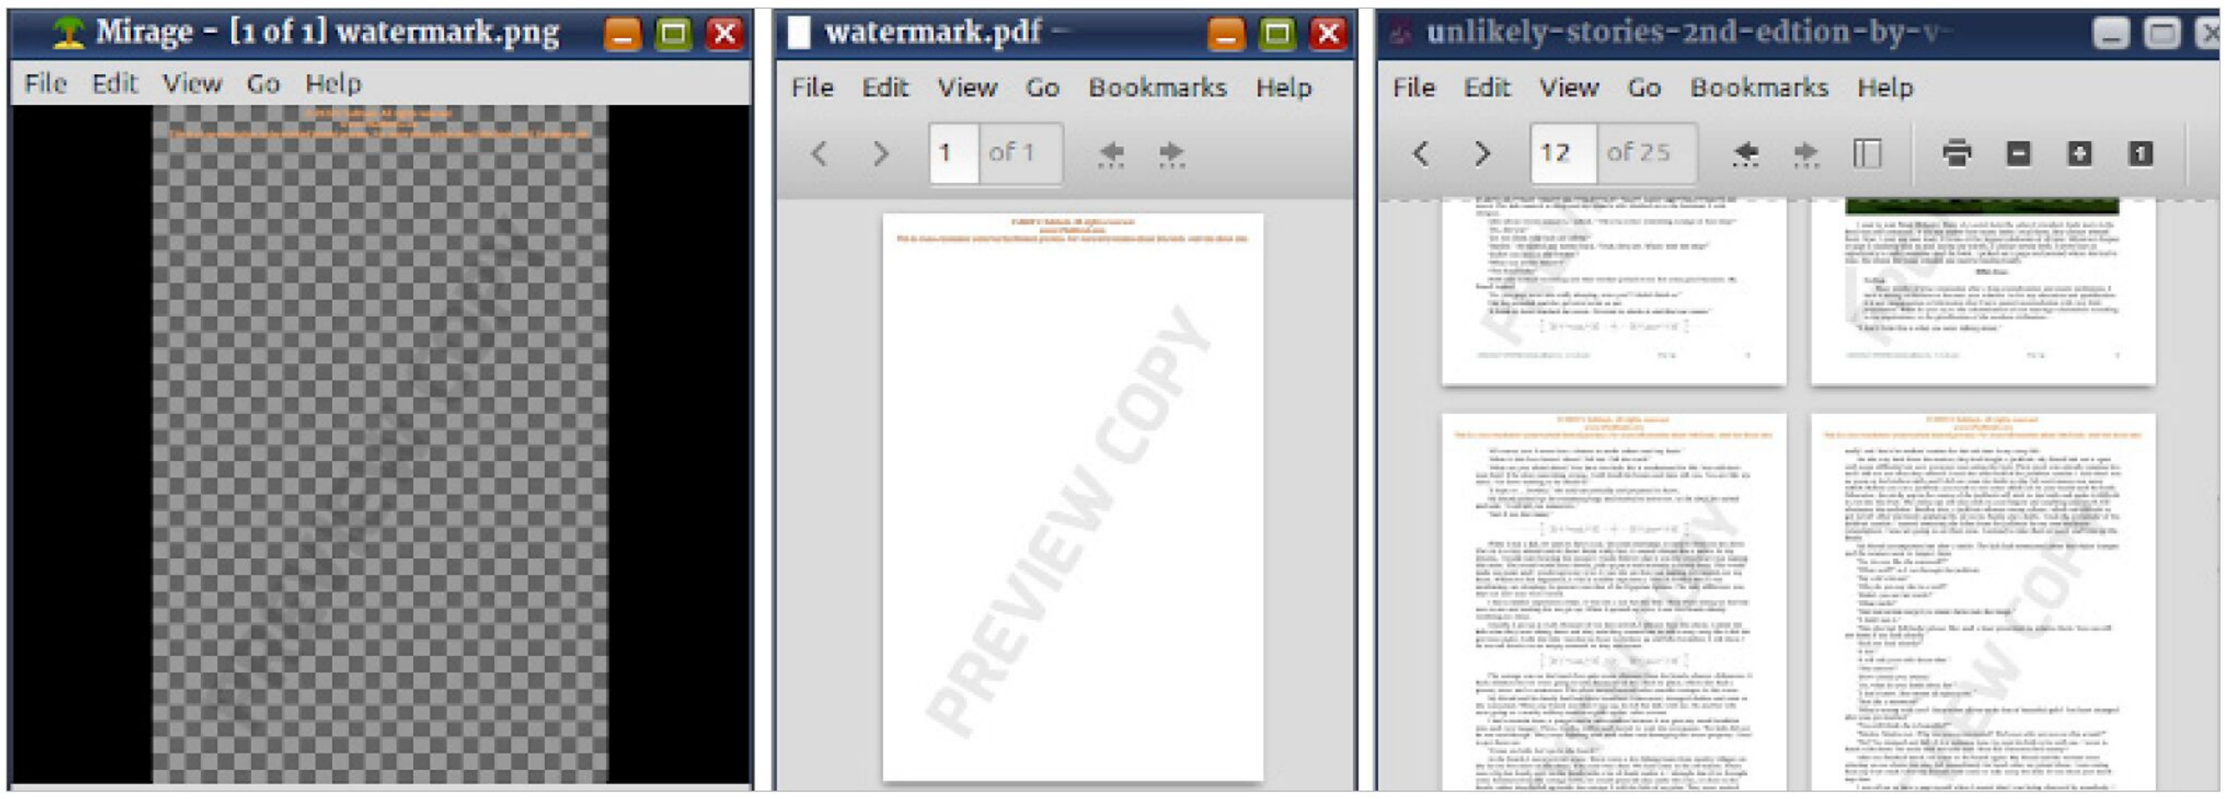


Figure 1: To stamp or watermark a PDF, first create a translucent image, convert it to PDF and then use it with *pdftk*

Splitting PDFs

There is no 'pdftk' option to remove pages. However, you can use its 'cat' option and specify pages that need to be left in the output document.

```
pdftk book.pdf cat 6-end \
      output story.pdf
# eliminates pages 1 to 5
```

pdftk does have a *burst* option to convert each page into a separate PDF.

```
pdftk book.pdf burst \
      output page%02d.pdf
```

Specify the format mask of the page number in the output PDF name, similar to how you display numbers with the *printf* function of C standard library.

Other tasks

pdftk can do a lot of other tasks. Just consult its help output. While it can add and extract file attachments to and from a PDF, respectively, do remember that if you send a PDF with a file attachment by email, mail servers will pre-emptively remove your mail from the inbox of the recipient. END

By: V. Subhash

The author has written 30 books including 'Quick Start Guide to FFmpeg (Apress)', 'Linux Command-Line Tips & Tricks', 'How To Install Solar' and 'Cool Electronic Projects'.

A Quick Look at New Age Trends in AI

As AI evolves, we must navigate the ethical and practical challenges it poses. Here's a summary of key AI trends and their significance in shaping the future of industries.



In the field of artificial intelligence, there is a common misconception that artificial intelligence (AI), machine learning (ML), deep learning (DL), genetic algorithms and neural networks are synonymous or do the same function but are merely matured by algorithms.

Actually, artificial intelligence (AI) is the traditional model devised to mimic human brain functions by computers using a technology-driven approach and mathematical and statistical models (e.g., probability). Later, machine learning (ML) was conceived to make programming models of brain functions and improve the behaviour of AI models through various learning approaches, storing the results of machine learning models for future predictions. This was inspired by the use of the neural network, which is the core storage layer of the human brain. To improve the processing efficiency of this neural network-based solution, a distributed model was implemented for faster processing and simulation of various brain functions like object detection and analytical abilities.

Genetic algorithms improve deep learning models with self-learning ability and intuitiveness to maximise the application of AIML models, reduce the effort in training them, and improve the knowledge base in neural networks. Generative AI, for example, uses large language models (LLM) to handle activities like content summarisation, content regeneration, and content extraction through self-learning capabilities. All these methods contribute to improving the ability of computer-based applications to emulate human-like thinking, reducing the effort in building and training the models, and cutting costs in time-to-market, with the flexibility to apply them in multi-domain solutions.

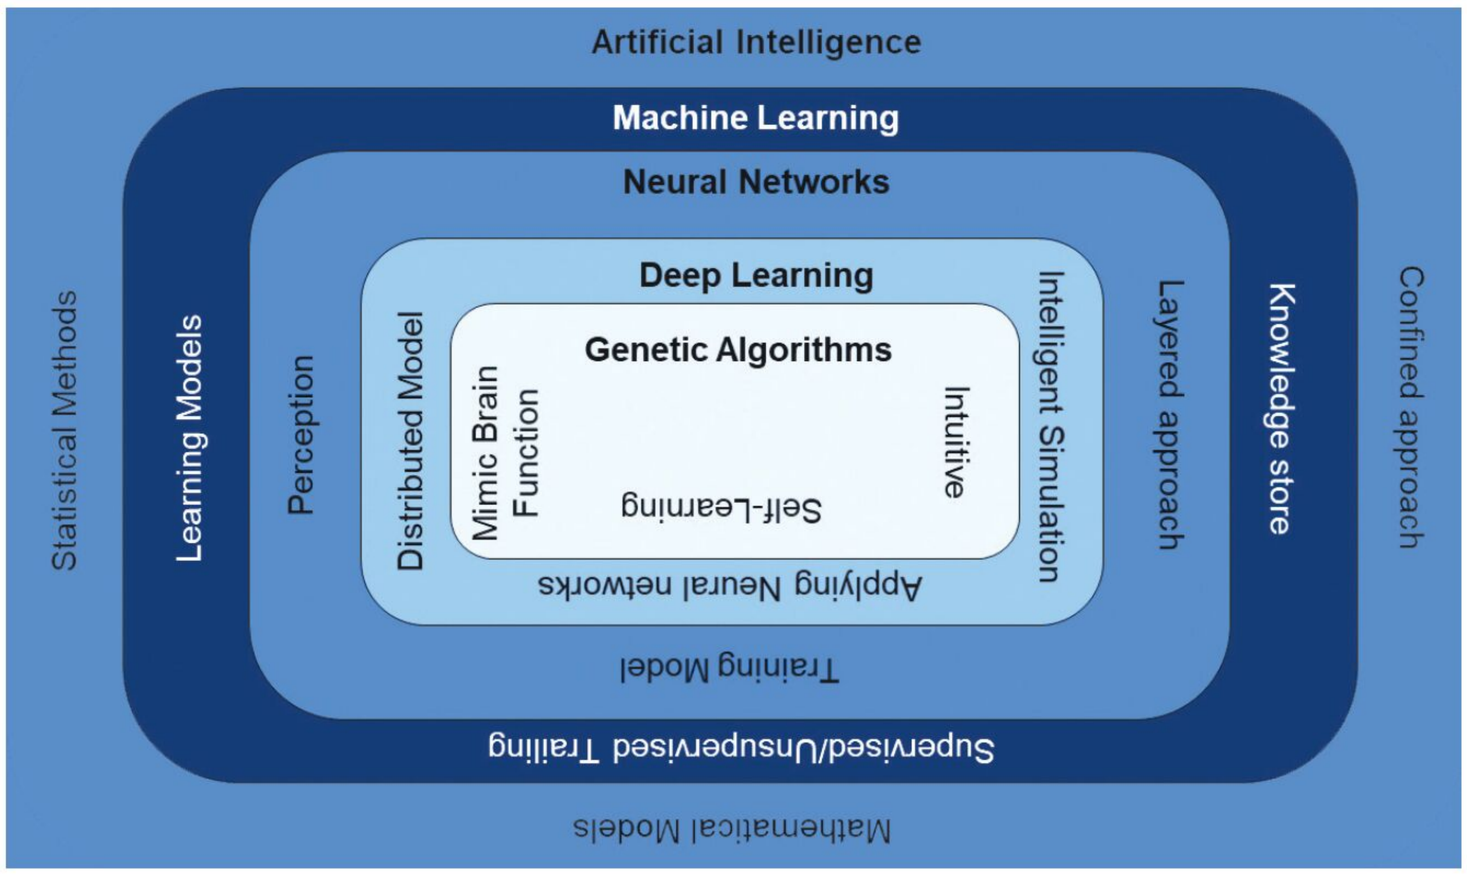


Figure 1: The paradigm of AIML

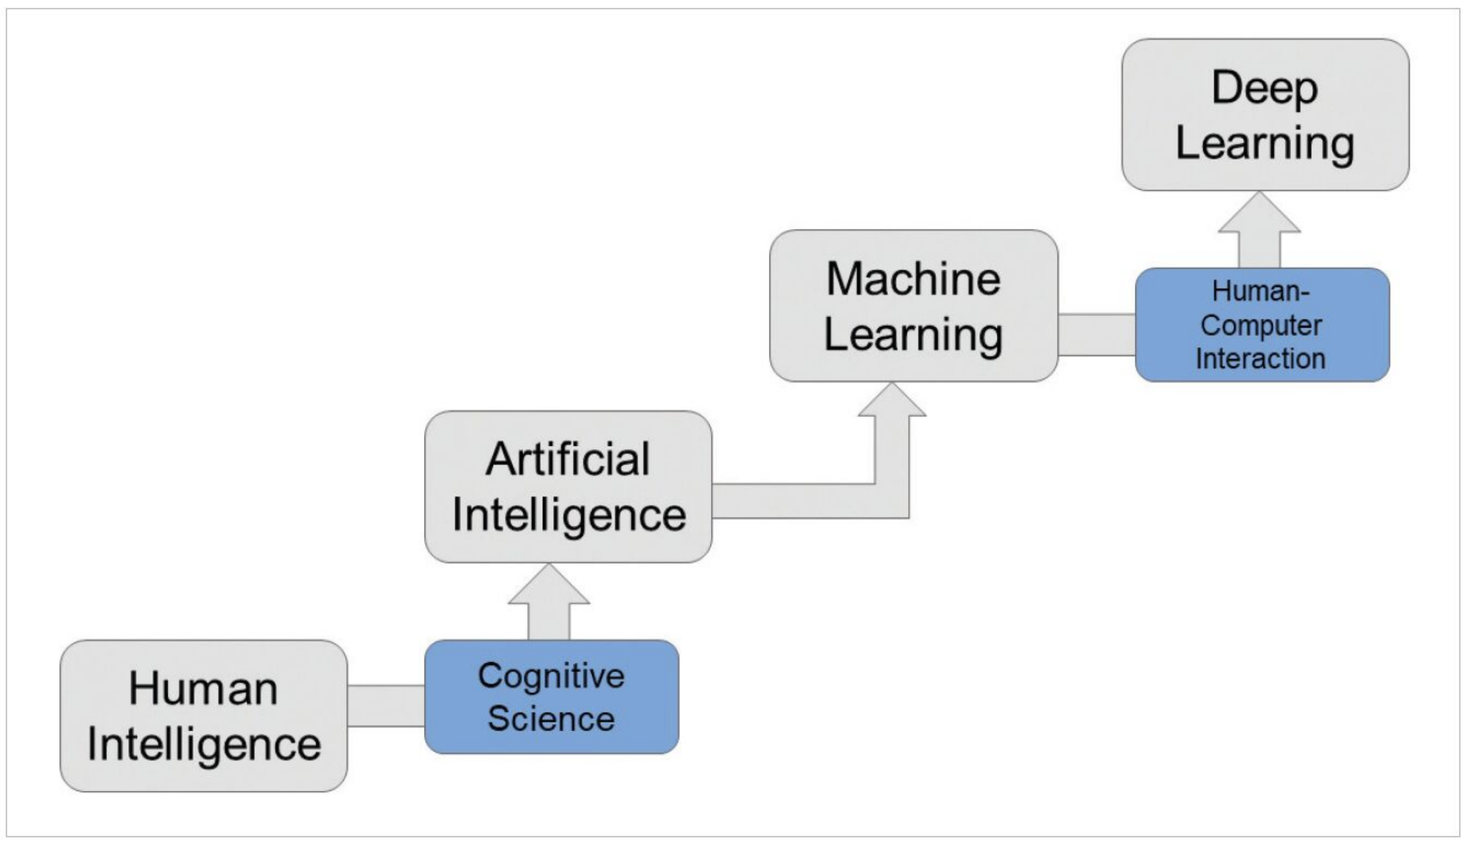


Figure 2: Singularity of human intelligence to deep learning

Importance of tracking AI trends

The importance of tracking AI trends can be summarised in the following points.

Stay ahead of the competition: AI is rapidly evolving, and businesses aiming to remain competitive need to be aware of the latest trends. By tracking AI trends, businesses can identify new opportunities and threats, and develop strategies to stay ahead.

Identify new market opportunities: AI is being used in a wide variety of industries, with new applications emerging all the time. Tracking AI trends helps businesses identify untapped market opportunities.

Make better strategic decisions: AI can be used to collect and analyse data, identify patterns, and make predictions. Businesses can use this information to make better strategic decisions about their products, services, and operations.

Improve efficiency and productivity: AI can automate tasks, improve decision-making, and optimise processes. By keeping an eye on AI trends, businesses can find ways to improve their efficiency and productivity.

Reduce costs: AI can automate tasks currently performed by humans, resulting in cost savings. Tracking AI trends enables businesses to identify cost-saving opportunities.

Improve customer experience: AI can personalise customer experiences, provide better customer service, and predict customer needs. Tracking AI trends helps businesses improve their customer experience.

Explainable AI

Explainable artificial intelligence (XAI) refers to the development of AI systems that can provide understandable and transparent explanations for their decision-making processes. XAI is essential because traditional AI models, like deep neural networks, often operate as 'black boxes', making it challenging to comprehend why they reach specific conclusions or predictions.

XAI aims to bridge this gap by enabling humans to understand and trust AI systems. It involves techniques such as model interpretability, feature importance analysis, and rule-based systems to make AI more transparent. This transparency enhances accountability, aids in debugging and improving models, and ensures that AI systems adhere to ethical and regulatory standards.

Various approaches to XAI exist, with the best approach depending on the specific AI model and application.

Feature importance: Identifies the most important features for an AI model's decision-making

Counterfactual explanations: Shows how an AI model's prediction would change if input features were different

Local explanations: Explains an AI model's decision for a specific input example

Model introspection: Allows users to explore an AI model's internal workings

XAI is a rapidly growing field with ongoing research, and has the potential to make AI systems more transparent, trustworthy, and fair.

As an example, a hospital can use XAI to explain how a machine learning model is making decisions about cancer treatment. This can help doctors to understand the model's reasoning and

to make better treatment decisions. XAI can also be used to identify biases in the data or the model, which can help to improve the accuracy and fairness of the model's predictions.

Generative AI

Generative AI is at the forefront of creative innovation. Through neural networks, this technology enables the generation of content, art, music, and even human-like text. Artists, designers, and authors are using generative AI to create one-of-a-kind and inventive works. AI-powered technologies aid in the creation of unique designs in the field of design, while generative AI composes tunes that fascinate listeners in the field of music. However, ethical concerns arise due to its potential for generating deepfake content and misinformation.

In recent years, generative AI has made considerable advances, allowing it to create highly realistic and imaginative content across a wide range of areas. The invention of generative adversarial networks (GANs), a type of machine learning model, is one of the most breakthrough advances in generative AI. GANs are made up of two neural networks: a generator and a discriminator. The generator generates new data instances, while the discriminator validates them. These two networks compete in a process in which the generator wants to create data that is indistinguishable from genuine data and the discriminator aims to correctly categorise real and generated data. The following are some of the most recent advances in Generative AI.

High-quality image synthesis: GANs can generate high-resolution images that often resemble real photographs. This can be used to create art, design, and even photorealistic representations of non-existent things or landscapes.

Style transfer and artistic rendering: Generative AI allows users to apply one image's style to another''s content. AI-generated art has emerged as a result, combining the qualities of many styles and materials.

Text-to-image synthesis: Generative AI can now create visuals that closely match the specified scenarios when given a textual cue, making it useful in sectors such as creative writing and advertising.

Video generation and manipulation: Advances in GANs simplify video generation and manipulation, enabling AI to produce new video content, alter current video footage, and even interpolate between frames to produce seamless transitions.

Voice and audio synthesis: Beyond images, Generative AI expands into audio synthesis, creating lifelike speech, music, and sound effects for entertainment, gaming, and voice assistants.

Drug discovery and molecule generation: Generative AI aids in identifying novel medicinal compounds by producing chemical structures that match specified criteria. This speeds up the drug discovery process by allowing researchers to investigate a broader range of possible molecules.

Realistic human face generation: AI models create incredibly convincing human faces, raising ethical issues about future applications in deepfake content creation.

AI-generated content curation: Generative AI supports content creators by suggesting ideas, generating drafts, and even composing music or text segments depending on specific topics or inputs.

Cross-domain applications: Generative AI finds applications across domains, allowing artists, designers, scientists, and engineers to use its creative powers in a variety of applications.

As generative AI advances, it has the potential to transform businesses, ignite creativity, and provide new ways to engage with and produce material. However, ethical concerns, such as the possibility of misuse and the need to ensure the responsible use of AI-generated content, must be addressed as this technology advances.

AI fairness and bias

Famous AI guru Ray Kurzweil coined the term 'singularity' in artificial intelligence, signifying the convergence of AI and human intelligence or making AIML accuracy approach human-like thinking. To achieve this highest level of accuracy in ML training, modelling, and functionality, it is of the utmost importance to address fairness and bias in AIML implementation.

Efforts to reduce or avoid bias in developing an AIML solution across various stages naturally lead to fairness in execution and results. Bias cannot occur on its own in the AIML stages by the system, but can result from human inputs.

For example, when preparing data for ML model training, one should have fairly distributed, unbiased data to prevent biased outcomes. Similarly, when labelling and grouping the data, training the model to simulate human-like thinking, deploying the model with knowledge gained from training, and interpreting the results from ML execution, it is essential to avoid pre-judgement or biased personal interests that could compromise fairness in building the entire system.

This includes avoiding race, income category, sexual orientation, gender, religion, marginalisation of the user and any discrimination of the user data when preparing the training data, as well as training the system and interpreting the results from the machine learning execution.

Various tools like IBM AI Fairness 360, Microsoft Fairlearn, and Google What-if are invaluable for validating and identifying if the training model and data collection has any bias in preparation and implementing fairness in the AI applications.

AI governance

AI governance is a set of principles, policies, and practices that organisations use to manage the development, deployment, and use of AI systems. It is important to have AI

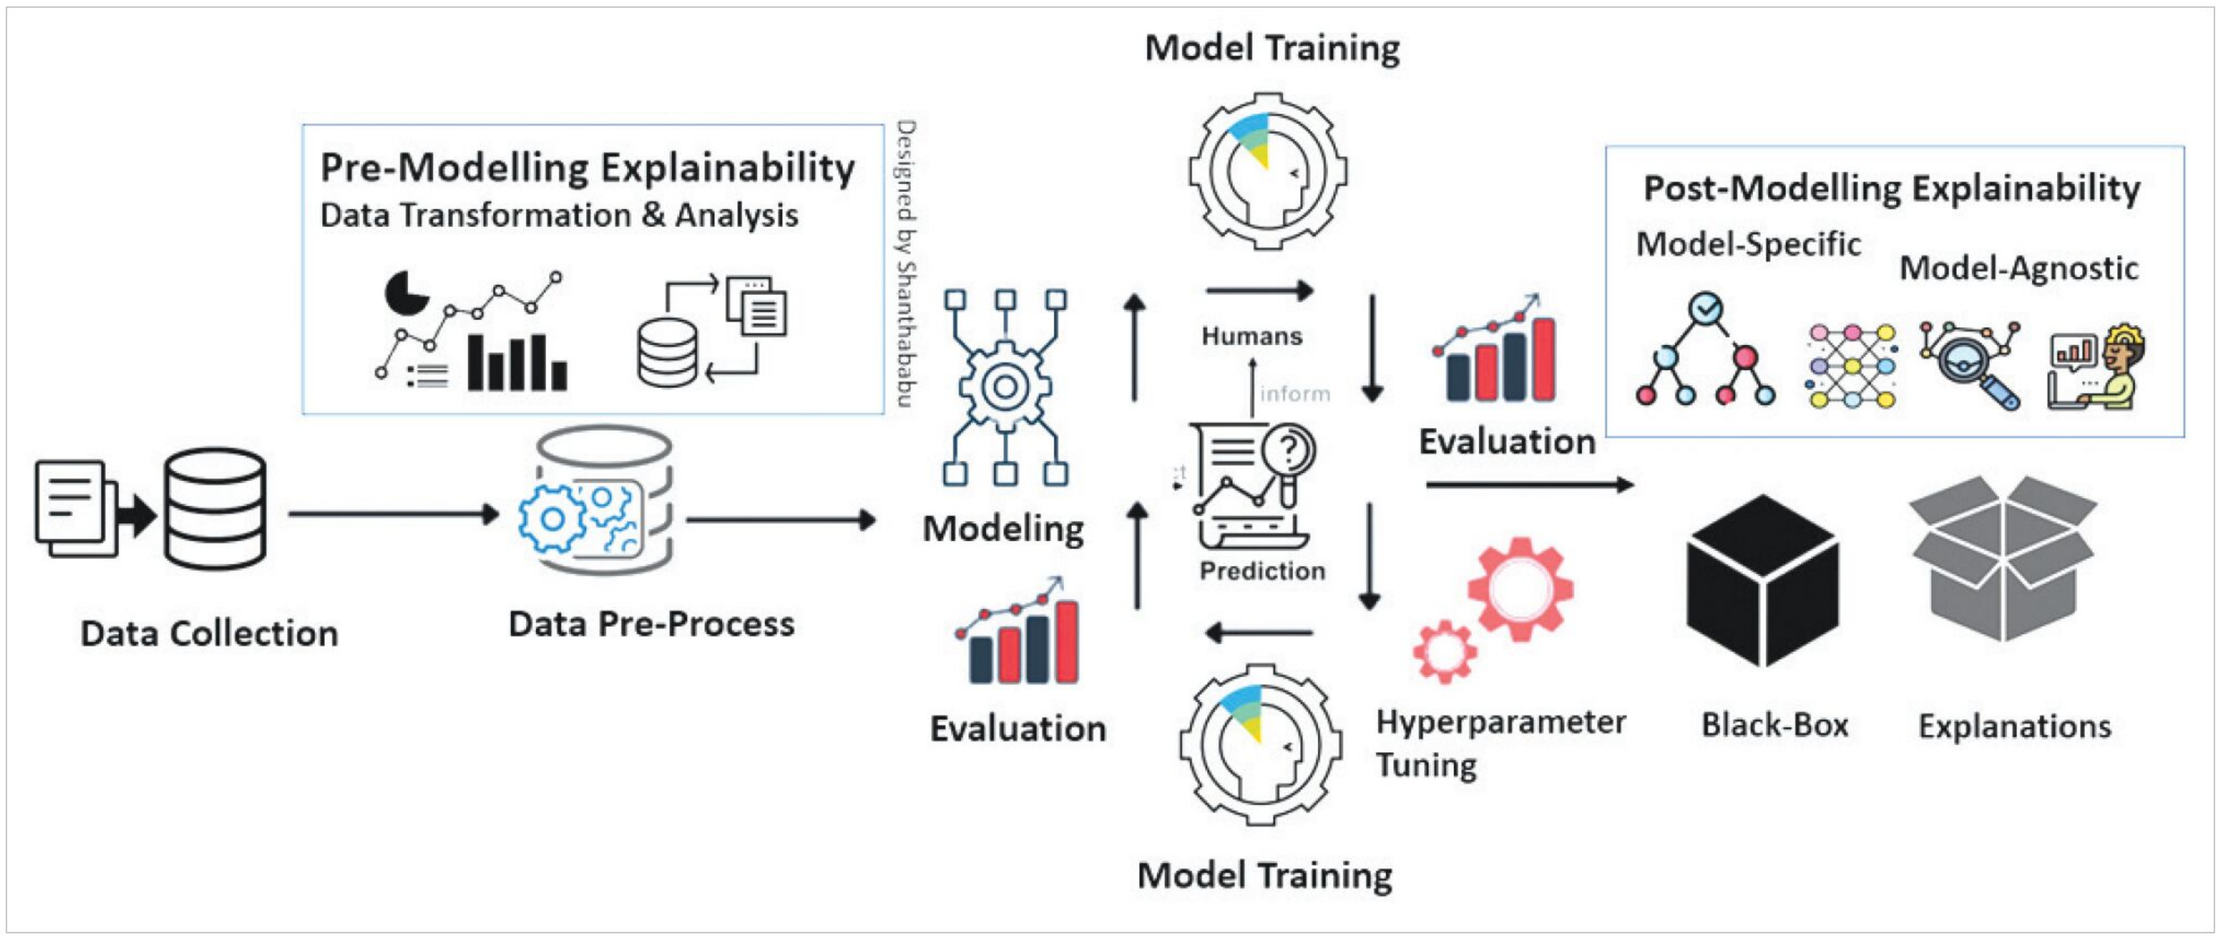


Figure 3: AI fairness and bias

governance in place to ensure that AI systems are used in a responsible and ethical way while mitigating the risks associated with AI.

Some of the key aspects of AI governance are:

Data governance: Managing data used for training and deploying AI systems to ensure that the data is accurate, complete, and unbiased.

Algorithmic accountability: Ensuring that AI systems are transparent and explainable. It is important to be able to understand how AI systems make decisions to identify and address any biases or unfair outcomes.

Human oversight: Ensuring that humans are involved in the development and deployment of AI systems so that these systems are used in a safe and responsible way.

Risk management: Identifying and mitigating the risks associated with AI systems including bias, discrimination, and job displacement.

Evolution of AI trends in the next five years

Artificial intelligence (AI) is a rapidly evolving field, and several trends are expected to shape the industry in the next five years. Here are some upcoming AI trends.

- ***Reinforcement learning and self-learning systems:*** These systems are expected to become more prevalent, capable of learning from their own experiences and improving their performance over time without human intervention.
- ***Autonomous vehicles:*** AI is a crucial component of autonomous vehicles, enabling them to perceive their surroundings and make decisions.
- ***Generative AI:*** The use of generative AI to create new content, such as images, videos, and music is expected to expand across various industries, including entertainment and advertising.
- ***Digital twinning:*** Creating digital replicas of physical objects or systems is set to become more prevalent in industries, including manufacturing and construction.
- ***Federated learning:*** This is a distributed machine learning approach that allows AI models to be trained on data that remains on users' devices. This can help to address privacy concerns and improve the security of AI systems.
- ***Multimodal learning:*** AI models can learn from multiple data types, such as text, images, and audio. This can help to improve AI system accuracy and robustness.

These trends are expected to have a significant impact on various industries, including healthcare, transportation, cybersecurity, and entertainment. As the field of AI evolves, these trends will shape the future of many industries. END

References

- *https://sitn.hms.harvard.edu/uncategorized/2020/fairness-machine-learning/*
- *https://www.ibm.com/analytics/common/smartpapers/ai-governance-smartpaper/#ai-governance-delivers*
- *https://www.synthesia.io/post/ai-trends*

By: Dr Magesh Kasthuri and Dr Anand Nayyar

Dr Magesh Kasthuri is a senior distinguished member of the technical staff and principal consultant at Wipro Ltd. This article expresses his views and not that of Wipro.

Dr Anand Nayyar is a PhD in wireless sensor networks and swarm intelligence. He works at Duy Tan University, Vietnam, and loves to explore open source technologies, IoT, cloud computing, deep learning and cyber security.

Open Source India 2023: WHERE THE INDUSTRY MET THE COMMUNITY

Here's a pictorial overview of all that took place last month at Open Source India, the mega event for open source enthusiasts held at Bengaluru.

October 2023 witnessed a mega open source event — the 20th edition of the annual, much-awaited Open Source India (OSI). Attended by a large number of open source enthusiasts from across India, the iconic convention was held on October 19 and 20, 2023 at the NIMHANS Convention Center in Bengaluru. It featured packed sessions by 80 speakers in six different conference tracks and 11 workshops — all running in parallel over two days.

A detailed hands-on workshop on 'Accelerating Generative AI Performance', which introduced Intel's OpenVINO as a solution, was a highlight of the event.

The keynote sessions in the conferences were delivered by experts from Amazon Web Services, Clari and GitHub. Day 2 witnessed keynote sessions by luminaries like Prof. Sowmya Narayanan Sadagopan from IIIT-Bangalore and Lothar Karl Becker from The Document Foundation. The afternoon keynote was delivered by Sanil Kumar D. from the SODA Foundation on 'Discovering Unstructured Data!'. The community also enjoyed the panel discussion on how 'Open Source Enables Startups'.

Apart from the talks by experts, the event was also an occasion for exhibitors to showcase their various community-centric offerings to thousands of attendees.

The audience at this mega 20th edition of OSI ranged from CIOs and tech heads to evangelists and academicians.

Here is a pictorial round-up for those who missed the event.

Powered By



1



2

Date: 19th-20th October 2023

Venue: NIMHANS Convention Centre, Bengaluru

Total registrations:

15,372

Total attendees:

4743

Total number of sessions:

46

Total number of workshops:

11

Total number of speakers:

80



3



4



5

01. Attendees visiting the Salesforce booth

02. Attendees visiting the Clari booth

03. Attendees visiting the PureStorage booth

04. Attendees at the registration counter

05. Attendees interacting with the exhibitors









The keynote sessions

06. Dr Biswajit Mohapatra from Amazon Web Services spoke on 'OpenSource AI/ML — AWS and the India Opportunity'
07. Sanil Kumar D. from SODA Foundation talked about 'Discovering Unstructured Data!'
08. Karan MV from GitHub delivered a keynote session on 'India Builds on Open Source'
09. Lothar Karl Becker from The Document Foundation spoke about LibreOffice and the awesome community behind it
10. Prof. Sowmya Narayanan Sadagopan expressed his views on how India is turning into an uploading nation
11. The keynote session on 'The True Power of Embeddings' by Srikar Yekollu, Clari







12. The panel discussion on 'Open Source Enables Startups' (left to right): Sunil Kalkunte, Code and Theory; Sanil Kumar D., Caze Labs; Divyanshu Verma, Redinent Innovations (Chair); Naveen Agrawal, Catastrophe Analytics; Abhas Abhinav, DeepRoot GNU/Linux; Alok Bardiya, WSO2

















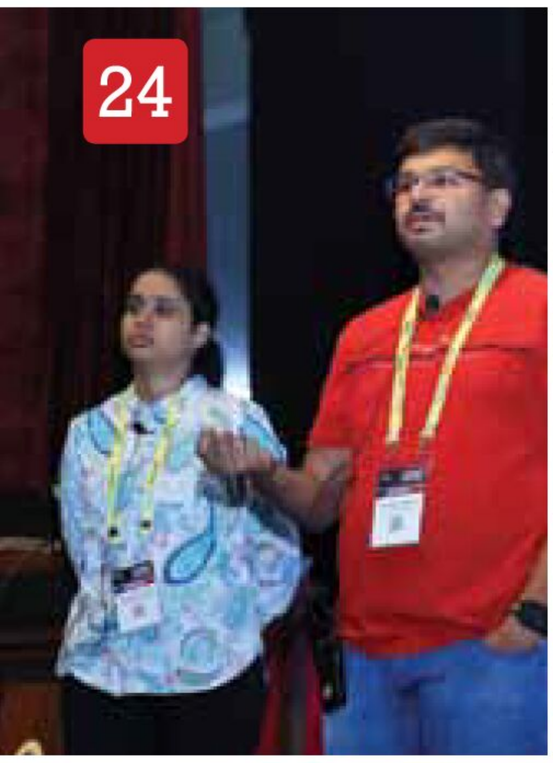






The workshops

13. Anisha Udayakumar from Intel conducted the workshop on 'From Concept to Code: Accelerating Generative AI Performance'
14. Mohan Sridharan and Ambily Kavumkal Kamalasanan, GitHub, interacted with attendees during their workshop on 'Building an App Alive with Generative AI'
15. Srikar Yekollu, Clari, conducted a workshop on 'A Hands-on Experience with Customising Embeddings'
16. 'Mastering Site 24x7: A Hands-On Monitoring Workshop' was conducted by Sarath Chandran from Zoho
17. Sreeja Radhakrishnan, Puja Chaudhury, Ramya Gogi, Karan Singh from Dell in their workshop on 'Managing a Stateful Application in K8s Using Go and Helm'
18. Mayank Raj, Salesforce.com India Pvt Ltd, conducted a workshop on 'GitOps Mastery: Declarative Configuration Management with Kubernetes'
19. Workshop on 'Transforming Business Processes using Open Source Low-Code Platform Joget' by Jeeva S. Chelladhurai, Comorin Consulting Services
20. Attendees learnt about 'Leveraging the Power of Open Source Headless CMS for Content Management' from OD10's Sangram Rath
21. Neependra Khare, CloudYuga Technologies Pvt Ltd, conducted a workshop on 'What Does the Tech Manager Need to Know About Kubernetes?'
22. 'Mastering Clean Architecture: Transforming Bad Code into Robust Systems'-- a workshop conducted by Rahul Garg from Thoughtworks Technologies
23. Amit Kumar, Sushantha Kumar, Uday Swamy and Sanil Kumar D. from Caze Labs Private Limited in their workshop on 'Let Us Do Kubernetes Monitoring Along with AI-powered Observability!'







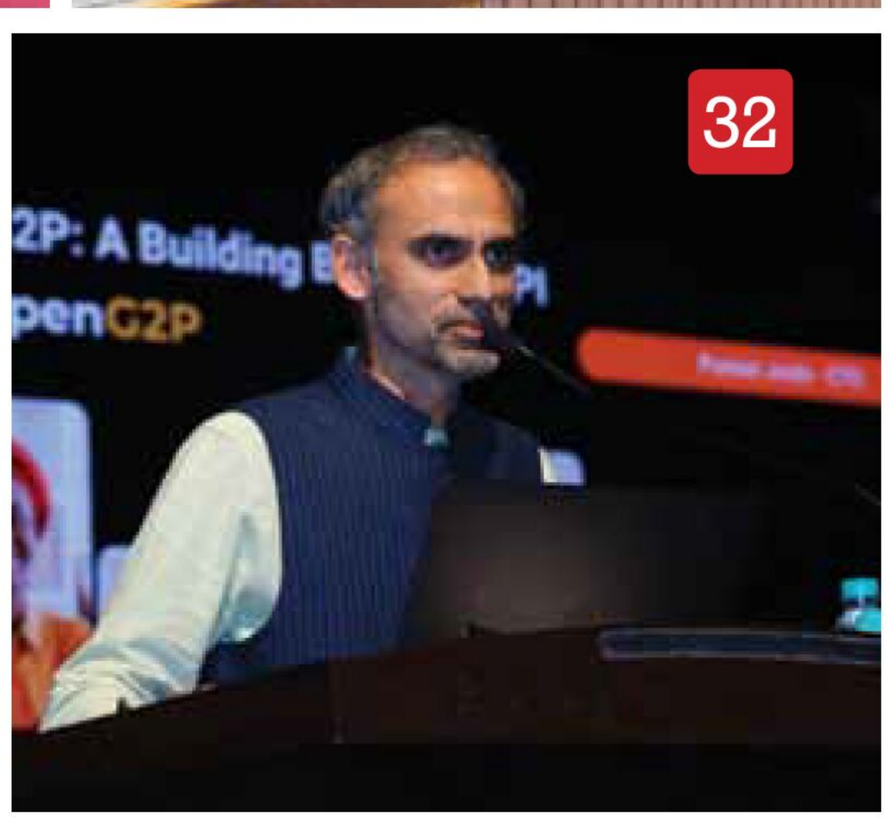














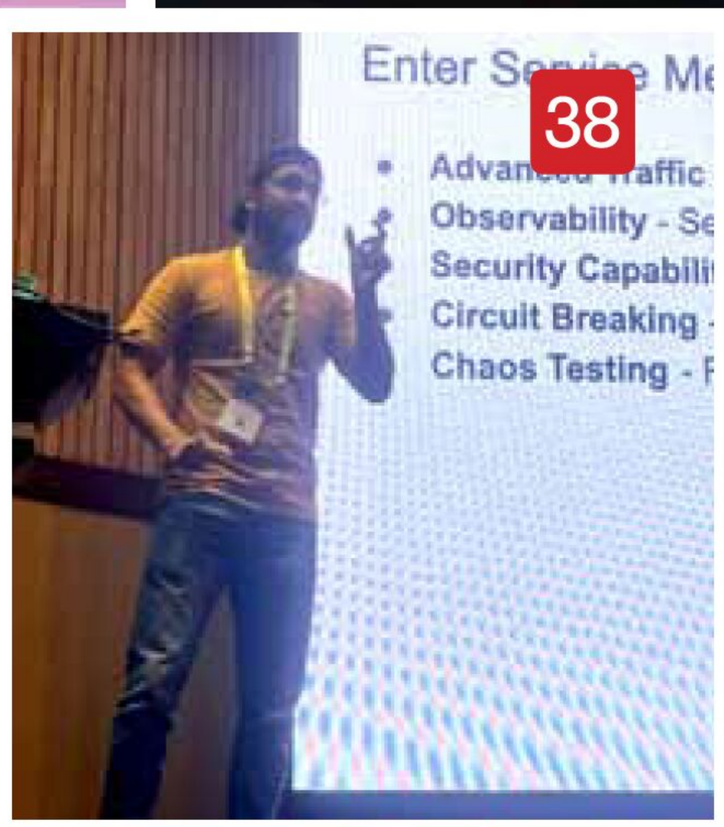
24. Procheta Bhattacharyya and Hariharan Thiagarajan from Thoughtworks spoke about 'Enabling the ONDC Revolution through Open Source Accelerators'
25. Ajeya Motaganahalli from Pure Storage shared his views on green data storage in the open source world
26. Laxminarayan Chandrashekar and Akash Sonune from Siemens Technology Services talked on 'Embracing Inner Sourcing Towards Open Sourcing'
27. Ananth Kumar from Zoho spoke on his journey with Cassandra
28. Preetdeep Kumar and Gonaiah Aare from Salesforce.com India Pvt Ltd in a session in the DevOps conference track
29. Vijay Kulari and Keerthivasan Suresh from Pure Storage sharing their knowledge
30. Arun Balaji and Aastha Varma, Amazon Web Services, talked about large language models
31. Amith K.K. from Unisys spoke on 'Building Open Source Generative AI Agents to Supercharge your Workflows'
32. Puneet Joshi from OpenG2P talked about 'OpenG2P: A Building Block for Digital Public Infrastructure'
33. Srivathsa N.S., Unisys, spoke on 'Generative AI and Intellectual Property (IP) - Understanding the Techno-legal Landscape'







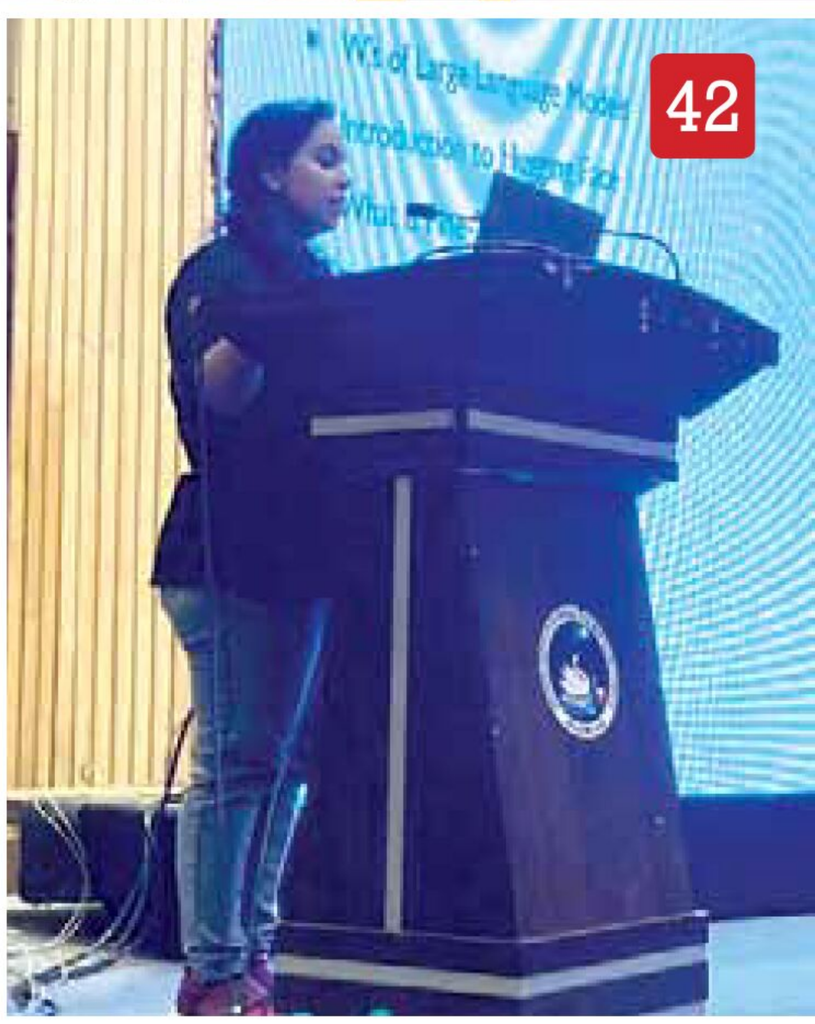












34. Biju K. Nair, Legalitech.In, gave his views on the legal landscape and Intellectual Property (IP)
35. Amit Upadhyay from fastn.com explained the new full stack programming language for web development
36. Zameer Fouzan, New Relic, spoke about OpenTelemetry
37. Leela Srinivas Siripurapu from Pure Storage talked on 'Reimagining MLOps with Kubeflow and Portworx'
38. Animesh Rastogi from Searce took a session on 'Progressive Delivery on Kubernetes'
39. Shubham Srivastava and Deepak Kumar from Zenduty presented a success story
40. Vivek Sridhar from Microsoft explained how to orchestrate your AI with Semantic Kernel
41. Sathyan Sethumadhavan and Karrtik Iyer from Thoughtworks Technologies talked about large language models
42. Binitha MT and Subhankar Adak from Dell Technologies spoke on 'Fine Tuning Large Language Mode' in the AI and ML conference track

















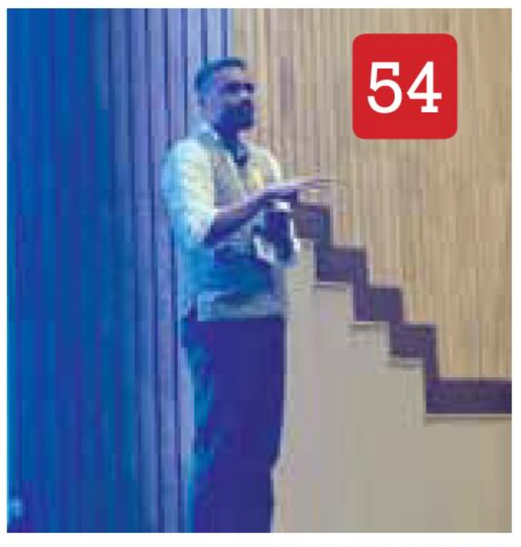


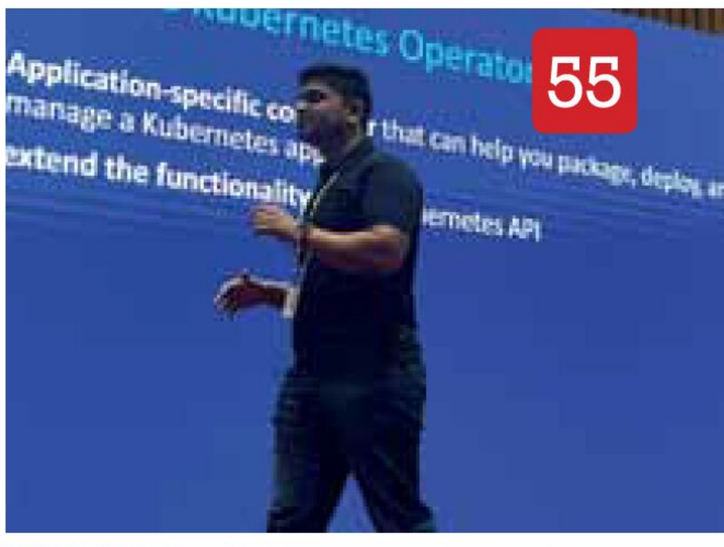


43. Ankit Sharma, Cisco Systems Inc., spoke on security for AI
44. Gopal S. Akshintala from Salesforce.com India Pvt Ltd elaborated on 'A Template-driven API Automation Tool for JVM'
45. Tony Wasserman, software and strategy advisor, gave his views on 'Organizational Adoption and Governance of Open Source Software'
46. Indumathi Venkataraman from Code and Theory spoke on 'How Open Source is Redefining Organisational Dynamics'
47. Sachin Garg, NavankurIT, in his session on 'Building a Community - Then and Now'
48. Akshita Agrawal, LegaliTech, spoke in the AI and ML conference track
49. Sourav Kumar Paul and Abhijit Kulkarni from Intel Corporation discussed the performance of Java
50. Sudipta Saha from Accenture spoke on DevSecOps in industrial data pipeline
51. Ankit Kumar from Aklivity spoke about 'Harnessing Zilla for Dynamic API Delivery in Event-Driven Architectures'



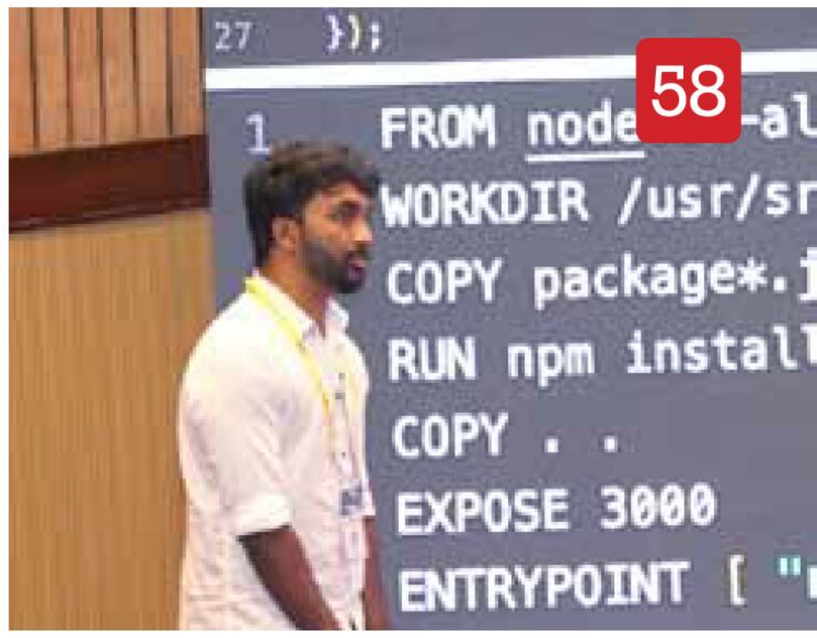








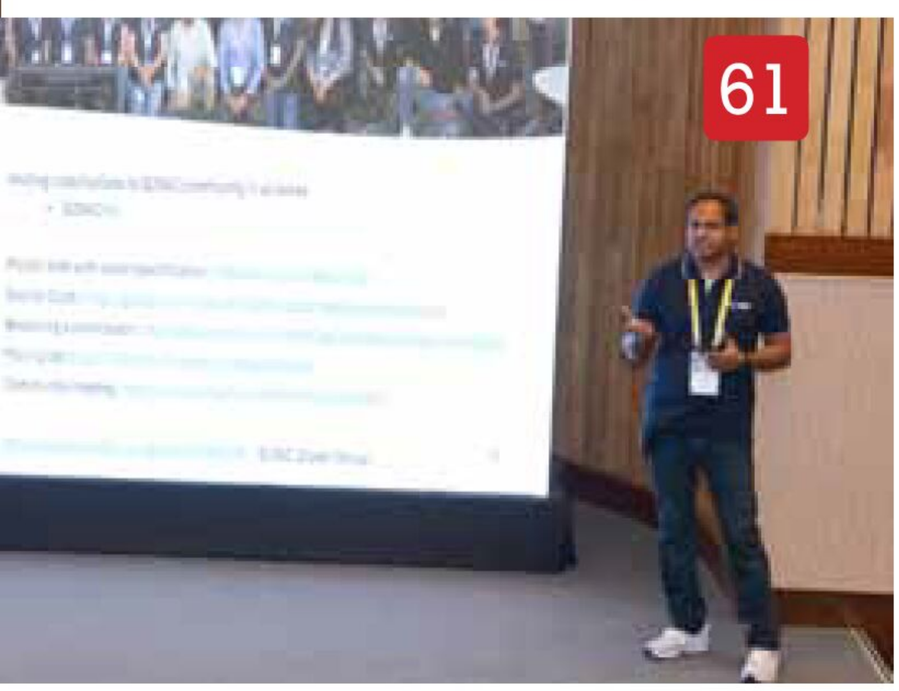


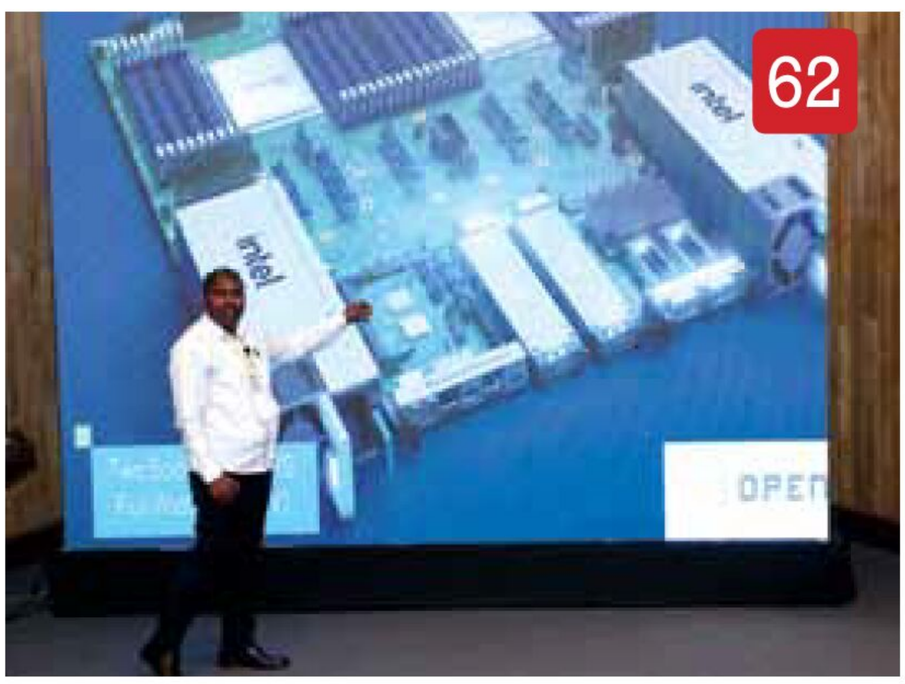


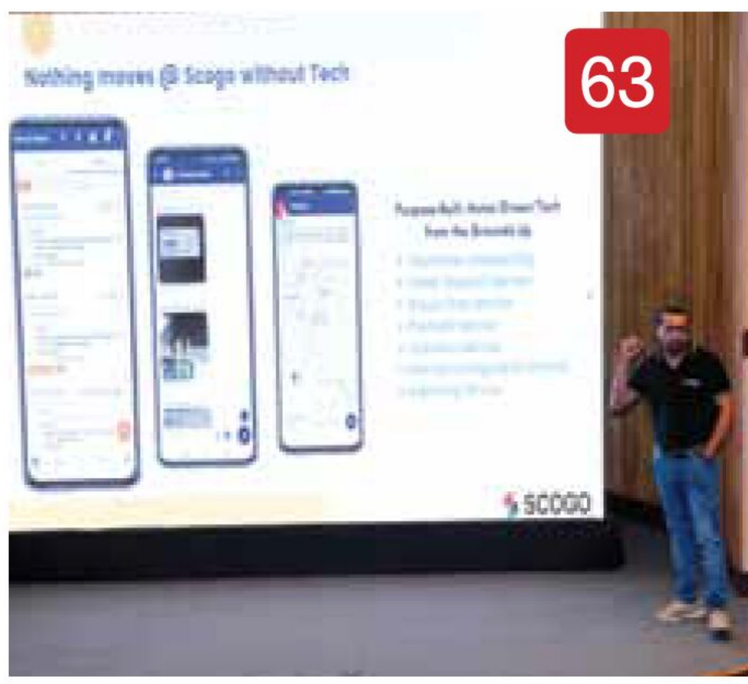








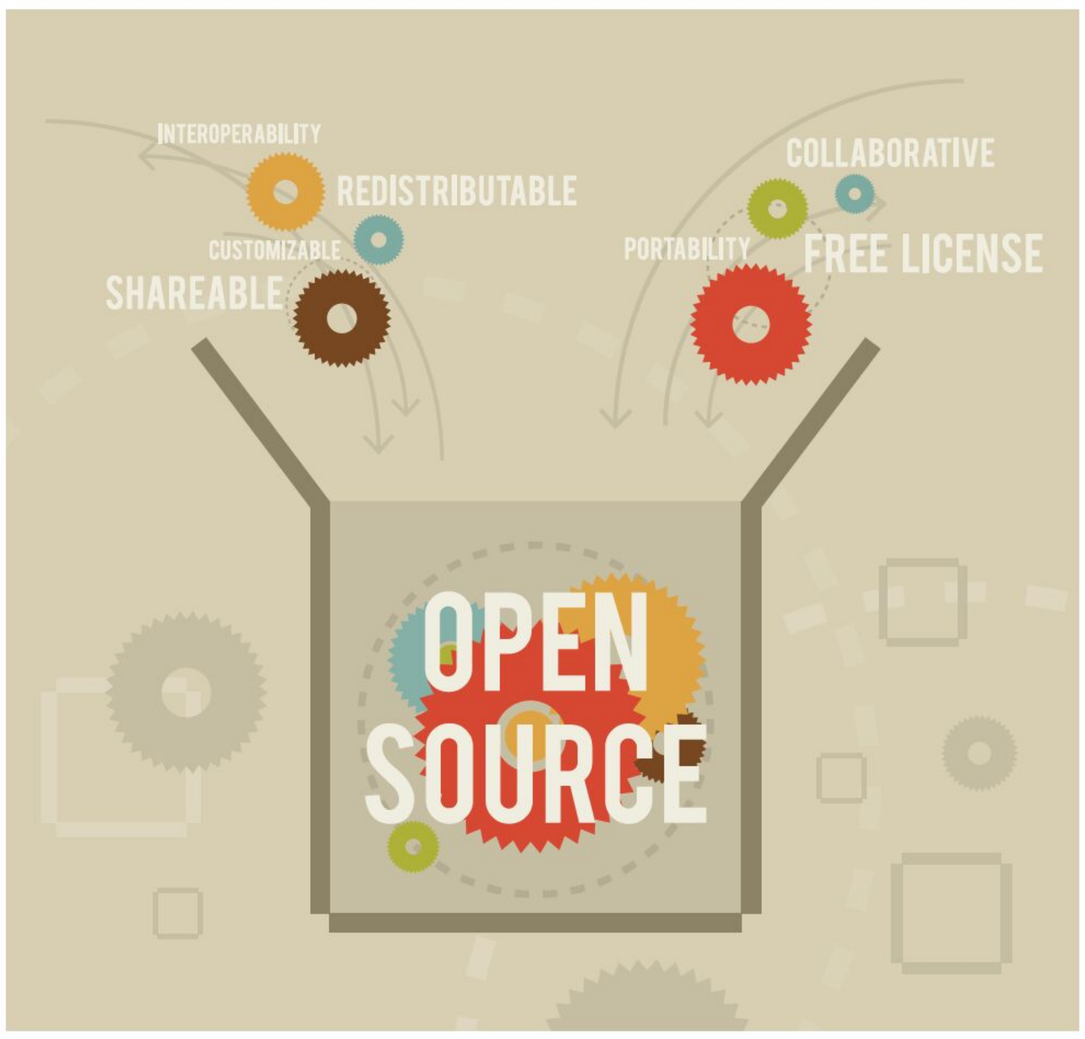




52. Shivji Kumar Jha from Nutanix spoke on 'OSS DB Architectures'
53. Vinoth Kanna from Mydbops talked about the MySQL ecosystem
54. Vijay P. from Tibil Solutions spoke on 'Unveiling Business Data Observability'
55. Aditya Joshi from Walmart shared his thoughts on 'Automating Workloads with Kubernetes Operator'
56. Abhishek Gupta from Amazon Web Services spoke about 'Writing Kubernetes Infrastructure as Go Code'
57. Vinay Rajagopal, Red Hat, spoke about the carbon footprint of our code and of containers
58. Goli Akshay Sujith from Slice took a session on NPM's impact on Dockerized Node.js applications
59. Baskaran from Dell spoke on 'Network Security in Kubernetes based Platforms'
60. Dharmanandana Reddy Pothula from Caze Labs spoke about cloud native security
61. Ashay Krishna from Microsoft spoke on 'SONiC - The Future of Networking OS'
62. Ravindra Reddy Padala from Dell spoke in the IT Infra conference track
63. Karan Singh from Scogo Networks Pvt Ltd during his session on 'Transforming the Tech Stack'

The Benefits and Drawbacks of OSS

Open source software has witnessed a steady increase in adoption over the past two decades, but the emergence of artificial intelligence, blockchain and edge computing has taken its popularity to an all-time high. While OSS has many benefits, it does come with a set of challenges.



A silent revolution is taking place in the technology and software development world, one that encourages collaboration, celebrates transparency, and empowers individuals and organisations. Open source software (OSS) is at the centre of this transformation, reshaping our interactions with technology and the digital universe.

The origins of OSS

OSS has its origins in the late 1970s and early 1980s, when the ethos of 'free software' began to take shape. Pioneers such as Richard Stallman, the founder of the Free Software Foundation, laid the groundwork for the OSS movement. Stallman's GNU Project, which began in 1983, aimed to create a free UNIX-like operating system that included all necessary software. This project was motivated by the belief that software should be freely shared, with users free to run, study, modify, and distribute it as they saw fit.

The term 'open source', as we know it today, became popular with the establishment of the Open Source Initiative (OSI) in 1998. The OSI sought to promote open source software adoption by emphasising the practical benefits of transparency and collaboration. It proposed the open source definition — a set of requirements for software to be considered open source. This definition included principles such as free redistribution, access to source code, and no discrimination against individuals or groups.

The first impact of open source was the liberation of software development from proprietary constraints. This novel approach to software development enabled individuals, organisations, and communities to collaborate on projects on an unprecedented scale. OSS projects sprouted up all over the digital landscape, laying the groundwork for a new era of technology.

The GNU/Linux operating system was one of the first and most influential OSS projects, a collaborative effort to create a free and open source UNIX-like operating system. Linus Torvalds launched Linux in the early 1990s, and it quickly gained traction as a dependable and versatile alternative to proprietary operating systems.

The success of Linux catalysed the adoption of open source practices in a variety of fields, inspiring a wide range of software projects to take an open approach. Apache, an open source web server, became the World Wide Web's foundation, serving many websites worldwide. The Mozilla project, responsible for the Firefox web browser, marked yet another milestone in spreading open source ideals. These developments demonstrated that not only was open source software feasible, but was also capable of outperforming its proprietary counterparts.

Why open source is important

- ***Collaboration and innovation:*** OSS is a hotbed of innovation. It promotes global developer collaboration, resulting in faster software development and enhancement. This collaborative approach enables businesses to tap into a vast pool of expertise, which frequently results in cutting-edge features and functionalities.
- ***Cost-efficiency:*** One of the primary benefits of OSS is the ability to save money. The use of open source software can be a game changer for businesses, particularly startups

Stay **Connected.**
Stay **Informed.**
Stay **Ahead.**

EFY GROUP
YOURS SINCE 1969



SUBSCRIBE AND **SAVE**

TEAR OUT

ORDER FORM

TEAR OUT

PRINT MAGAZINE	1 YEAR (12 copies each)	3 YEARS (36 copies each)	5 YEARS (60 copies each)
Electronics For You (Rs 100/copy)		**WITHIN INDIA (IN RUPEES)**	
	840 ☐	2150 ☐	3000 ☐
		SAARC COUNTRIES (IN US$ BY AIR)	
	50 ☐	135 ☐	NA
		OTHER COUNTRIES (IN US$ BY AIR)	
	100 ☐	270 ☐	NA

e-magazine subscriptions within India are available at half the rates mentioned here.

Overseas rates for each e-magazine in US$: 1 year: $12; 3 years: $33; 5 years: $50 only

PRINT MAGAZINE SUBSCRIBERS GET:

- Free e-magazine every month
- Free delivery of print magazine by post
- And much more (check: subscribe@efy.in)
- For delivery by courier, please add Rs 50 for each copy

To subscribe online, visit https://payment.efyindia.com | **OR** | **SCAN THIS CODE**

Name_____Organisation_____Mailing Address_____

_____City_____

Pin Code_____State_____Phone No._____ Email_____

Subscription No. (for existing subscribers only)_____. I **would like to subscribe to the above (✓)marked magazine(s) starting with the next issue.** Please find enclosed a sum of Rs_____by DD/MO/crossed cheque bearing the No._____ dt._____in favour of EFY Enterprises Pvt Ltd, payable at Delhi.

Please mark one (nearest) relating to your subscription: ☐ **Indian Company** ☐ **MNC** ☐ **R&D organisation** ☐ **Engineering institute** ☐ **College/School** ☐ **Any other (specify):**_____

Send this filled-in form or copy to : EFY Enterprises Pvt Ltd, D-87/1 Okhla Industrial Area, Phase 1, New Delhi 110 020 | Ph: 011-40596600 | e-mail: support@efy.in

Terms:- # These rates are applicable for new subscribers as well as renewal by existing subscribers # Please allow 4-6 weeks for processing of your subscription.
Please include your pincode for prompt delivery of your copy.

with limited resources. It eliminates the need for expensive licensing fees, allowing businesses to redirect funds to more critical areas such as product development and marketing.

- ***Customisation and flexibility:*** OSS allows businesses to customise software to meet specific needs. Startups, in particular, can tailor their software solutions to match their business models and customer expectations.
- ***Security and dependability:*** OSS's transparency improves security. Vulnerabilities are quickly identified and addressed when the source code is available for inspection. As a result, open source software is a dependable choice for corporations and startups alike, ensuring the stability of critical applications.

The impact of OSS on corporations

- ***Cost savings:*** Corporations have embraced OSS to cut operational costs. The elimination of licensing fees, combined with the ability to customise software to meet specific requirements, results in significant cost savings. Hence, OSS has become an appealing option for large corporations looking to optimise their budgets.
- ***Scalability:*** To meet growing demands, large corporations frequently require scalable solutions. Without the constraints of proprietary licensing, OSS provides the scalability required to adapt to changing circumstances. It ensures that companies can expand without incurring ever-increasing software costs.
- ***Project contributions:*** Many corporations actively contribute to open source projects. They can thus influence the direction and features of the software on which they rely. This level of participation ensures that the software meets their specific business requirements.
- ***Used in critical functions:*** Open source software is no longer limited to small projects; it now includes enterprise-grade solutions. Corporations can use robust open source solutions for critical functions like customer relationship management, content management, and cloud infrastructure.

The impact on startups

- ***Cost-effective innovation:*** Every rupee counts for startups. OSS provides a means for low-cost innovation. Startups can build and scale their products using open source tools and frameworks without incurring high software acquisition costs.
- ***Rapid development:*** In order to gain competitive advantage, startups must move quickly. The use of open source software shortens development cycles. Pre-existing open source components and libraries enable startups to focus on developing unique features while decreasing time-to-market.
- ***Global assistance:*** Startups can benefit from the assistance of open source communities. These communities provide resources for problem solving, knowledge sharing, and guidance. It's like having a global team of experts on hand to provide critical assistance.
- ***Adaptability:*** Open source solutions are highly adaptable, allowing startups to pivot or add new features without being constrained by proprietary licences.

The OSS advantage

OSS has advanced in a variety of ways since its inception to its current status as a driving force in the tech world. Here are some benefits that have been developed and refined over time.

- ***A diverse and thriving ecosystem:*** The proliferation of open source software has resulted in a diverse and vibrant ecosystem. A plethora of projects, libraries, and tools serve a variety of purposes and industries. This variety enables developers and businesses to find solutions that precisely meet their requirements. There is almost certainly an open source solution for a robust content management system, an intricate machine learning framework, or a versatile e-commerce platform.
- ***Improved security procedures:*** Open source software's security practices have significantly improved over time. The open source code allows a global community of developers and security experts to examine and improve the software's security features. Vulnerabilities are identified and addressed more quickly, which contributes to overall security.
- ***Collaboration on a large scale:*** The global contributor community has developed practices and tools that enable effective and efficient collaboration. Version control systems such as Git and collaboration platforms such as GitHub have proven useful in coordinating contributions from developers all over the world. This level of collaboration ensures that projects are developed quickly and collective wisdom is shared.
- ***Extensive documentation:*** Adoption and use of open source software require extensive and easily accessible documentation. Projects have invested in creating detailed documentation in response to this need. Users can find guides, tutorials, and reference materials to help them understand and effectively use the software. The availability of comprehensive documentation facilitates onboarding and encourages a larger user base.
- ***Industry-wide versatility:*** Open source software has proven to be extremely versatile in a variety of industries, including healthcare, finance, education, and e-commerce. As industries become more reliant on technology, open source solutions can be tailored to meet the unique needs of each sector, fostering innovation and efficiency.
- ***Web technologies:*** The growth of the internet has been supported by open source technologies. The web's backbone consists of technologies such as the Apache web server, the MySQL database, and the Linux operating

system. They have enabled the development and upkeep of numerous websites and web applications. This pervasiveness exemplifies the impact that open source has had on global connectivity.

- ***Adoption by corporations:*** Large corporations have recognised the value of open source software and implemented it in their operations. Many have established open source offices and become active contributors to open source projects.
- ***Impact on emerging technologies:*** Open source software has shaped emerging technologies such as artificial intelligence, blockchain, and edge computing. The availability of open source frameworks and libraries has accelerated the development of these technologies and made them more accessible to a wider range of developers.
- ***Educational institution accessibility:*** Open source software has been adopted by educational institutions ranging from primary schools to universities. Its low cost, adaptability, and educational licences make it an excellent choice for students and educators. This ensures that the next generation of developers understand the principles and practices of open source.
- ***Improved user experience:*** Open source projects have made investments to improve the user experience. Many open source applications now have user friendly interfaces, intuitive navigation, and user-centric design. This emphasis on user experience has increased the appeal of open source software to a broader audience.
- ***Better licensing models:*** The evolution of various licensing models has provided open source projects with greater flexibility and clarity. Licences such as the GNU General Public License (GPL), the MIT License, and the Apache License provide different options for sharing and modifying code. These licences have contributed to a better understanding of the legal landscape surrounding open source software.
- ***Strong ecosystem of support:*** A robust support ecosystem is often beneficial to open source software. Companies, communities, and individuals provide support and consulting services to ensure that users have access to the help they require. This infrastructure improves the dependability of open source solutions.

Open source software is a dynamic and enduring force in the world of technology due to its diverse ecosystem, commitment to security, global collaboration, and adaptability to various industries. However, the open source journey has not been without its set of challenges.

Drawbacks of OSS

- ***Lack of structured support*:** One of the primary disadvantages of OSS is the lack of structured, guaranteed support. Users frequently rely on community forums, user documentation, or volunteer help. While this approach has the potential to be extremely effective, it may not be appropriate for organisations or individuals with mission-critical applications. Downtime and operational challenges can result from a lack of timely and consistent support.
- ***Problems with compatibility:*** Integrating various open source components can lead to compatibility issues. Different versions of software libraries may be used in different projects, and certain configurations may be required. It can be difficult to ensure that these components work well together, especially when combining multiple open source tools.
- ***Licensing complexity:*** Open source licences can be complex and varied. This variety can cause confusion and legal issues for users who do not fully understand the implications of different licences. The possibility of inadvertently violating licence terms or obligations is a significant disadvantage.
- ***Limited documentation:*** While open source projects have made strides in providing comprehensive documentation, some still lack clear and user-friendly instructions. Inadequate documentation can prevent users from fully understanding and utilising the features of the software.
- ***Learning curve:*** Open source software can have a steep learning curve for newcomers. This disadvantage may deter adoption and limit the user base. This challenge may be exacerbated by a lack of user-friendly interfaces or intuitive designs.

What modern trends should focus on

Modern trends in OSS may inadvertently overlook several critical aspects that can impact the ecosystem's growth and success.

- ***Commercial viability:*** Some open source projects struggle to find a viable business model, potentially limiting their long-term viability and discouraging contributors from investing their time and effort. Open source initiatives need to focus on fostering commercial success.
- ***Accessibility:*** Open source software is widely available, but it might not reach underserved communities or those with limited internet access. The aim should be to bridge the digital divide and make open source solutions more widely accessible.
- ***Design and usability:*** Many open source projects prioritise functionality over design and usability. Modern design trends should emphasise user-centric design, improving user experience (UX) and user interface (UI), which can broaden open source software's appeal.
- ***Marketplace integration:*** Contemporary trends should explore ways to incorporate open source software into larger marketplaces and ecosystems. This integration can make it easier for a broader user base to discover and adopt open source solutions.

Proposed improvements

- ***Structured support models:*** Exploring structured support models can benefit open source projects. These models could include the establishment of commercial entities that provide open source software support, training, and consulting services. These efforts can provide users with dependable assistance and create long-term business opportunities.
- ***Improved licence education:*** Open source projects can improve education and awareness efforts to address licensing issues. Providing clear, accessible information about different licences and their implications helps users make informed decisions.
- ***User-centred design:*** Emphasising user-centred design within open source software can enhance its appeal and user-friendliness. Encouraging contributors to prioritise UX and UI improvements can help to improve overall user experience.
- ***Initiatives to improve accessibility:*** Open source projects and communities should actively participate in initiatives to improve accessibility. Collaboration with organisations and governments can expand digital inclusion efforts.
- ***Development of business ecosystems:*** Encouraging the development of business ecosystems centred on open source software can boost its commercial viability. Support services, paid add-ons, and value-added products can generate revenue for contributors while benefiting users.

Security threats and recovery in OSS

Security threats in the world of open source can be daunting, but the community driven nature of these projects often enables quick responses and effective recovery efforts. This section explores some notable security threats and how they were dealt with.

- ***Heartbleed (2014):*** Heartbleed, a well-known security flaw in open source software, was caused by a flaw in the OpenSSL cryptography library. It gave hackers access to sensitive data on servers, such as user names and passwords. The open source community responded quickly by issuing a patch and urging users to update their software. Major technology companies and organisations also assisted with vulnerability assessments and patch deployment, which aided in the recovery effort.
- ***Shellshock (2014):*** This was a vulnerability in the Bash shell, a popular open source command-line interpreter. Hackers took advantage of this flaw to execute arbitrary commands on vulnerable systems. Immediate patching, the creation of firewall rules, and thorough security audits were all part of the recovery effort. The free and open source community worked together to patch the vulnerability and release updates for affected systems.
- ***Equifax data breach (2017):*** The Equifax data breach was caused by a flaw in the Apache Struts open source framework. Recovery efforts included patching the vulnerability, conducting a thorough investigation, and implementing enhanced security measures.
- ***Log4Shell vulnerability in Apache Log4j (2021):*** The Log4Shell vulnerability in Apache Log4j affected the widely used open source logging library Log4j. This critical flaw allowed remote code execution and posed a serious threat to many systems. The immediate release of patches, urging users to update, and enhanced monitoring were all part of the recovery efforts.
- ***WordPress plugin vulnerabilities:*** WordPress, a popular open source content management system, has faced several security threats due to plugin vulnerabilities. Typically, recovery entails patching or removing vulnerable plugins, updating the WordPress core, and enhancing security configurations. The WordPress community emphasises the importance of monitoring and responding to plugin vulnerabilities as soon as possible.
- ***Supply chain attacks:*** Supply chain attacks on open source components are becoming more common. Attackers compromise open source package development and distribution by injecting malicious code. Identifying and mitigating compromised components, improving security measures, and promoting code signing and verification are all part of the recovery effort.
- ***Vulnerabilities in Apache Struts:*** Apache Struts, an open source web application framework, has been plagued by several flaws over the years. Patching affected versions, conducting code reviews, and improving the framework's overall security are all part of the recovery process.
- ***Cross-Site Scripting (XSS) and SQL injection attacks:*** XSS and SQL injection attacks are common web application security threats that frequently target open source projects. Identifying and patching vulnerabilities, implementing secure coding practices, and educating developers and users about best practices for mitigating these threats are all part of the recovery process.

To put it simply, while open source software has come a long way, recognising and addressing its challenges is crucial for its continued growth and success. By developing structured support models, improving licence education, emphasising user-centred design, promoting accessibility, and fostering business ecosystems, the open source community can mitigate these challenges with greater impact. END

By: S.V. Rohith

The author specialises in RPA, Python and Web 3.0, and has over five years of expertise in creating and managing Telegram bots.

FOCUS

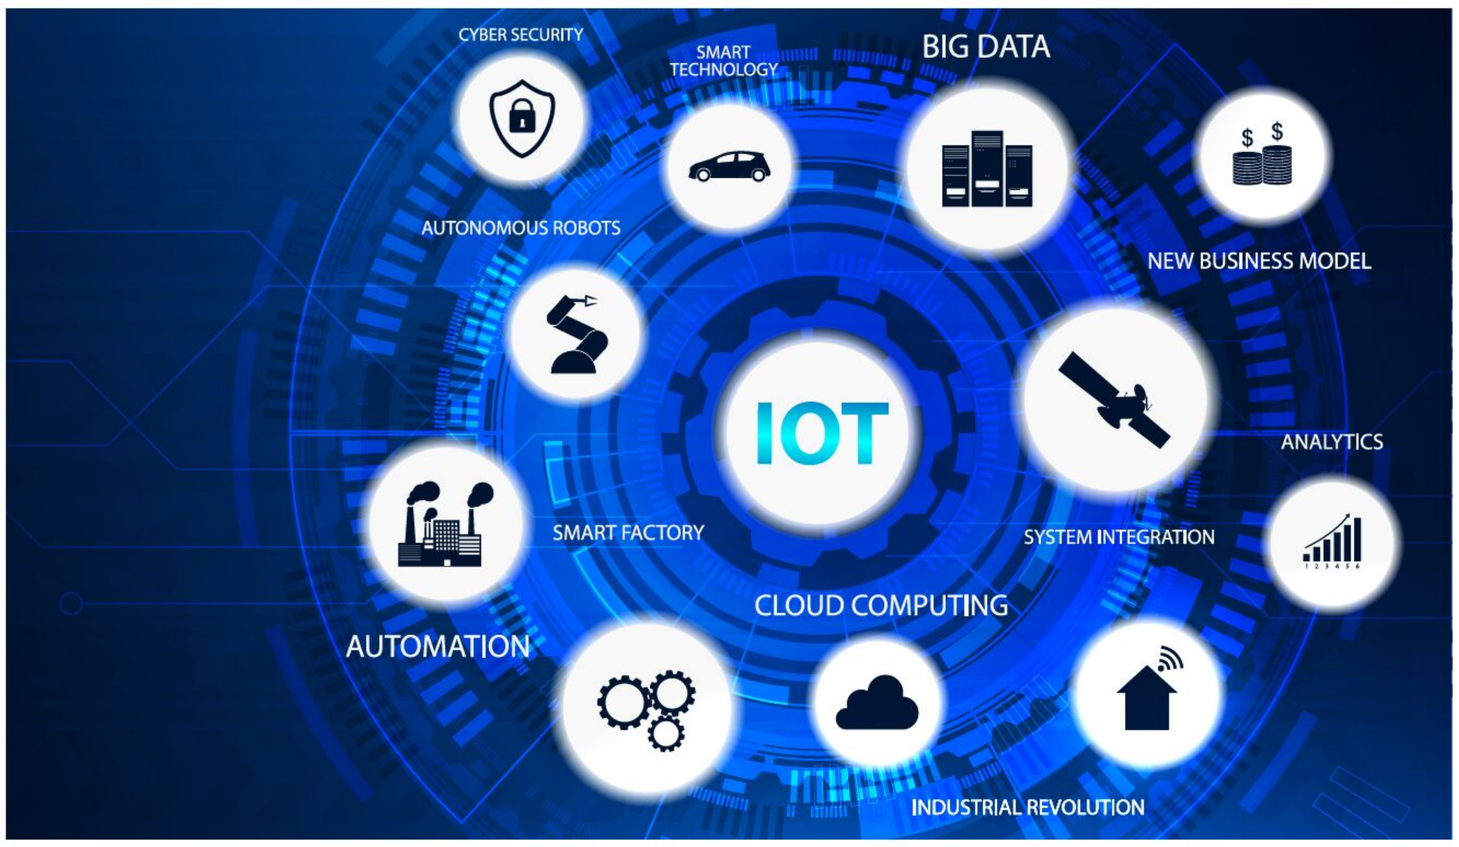
Internet of Things

Generative AI in IoT: Moving Towards Transformative Applications



Generative AI's contribution to the Internet of Things (IoT) cannot be underestimated. It has the potential to truly transform this domain by improving the security of IoT networks and making IoT devices more affordable, among other things.. *Read more on page.....40*

The Role of Open Source in the Growth and Adoption of IoT



The symbiotic relationship between open source and IoT is undeniable. Open source technologies have been instrumental in fuelling IoT's growth and widespread adoption. This relationship is underpinned by the principles of accessibility, flexibility, interoperability, security, and innovation that open source brings to the IoT landscape. *Read more on page.....43*

Neuromorphic Computing

Neuromorphic computing transcends mere connectivity, creating a future where our tech mirrors human cognition. Get ready to embrace brain-inspired tech!

Read more on page.....61

How Open Source is Helping IoT Evolve

Let's delve into the essence of IoT, explore its profound impact across diverse industries, and introduce the pivotal role of open source software in shaping its evolution.

Read more on page.....69

Leveraging the Power of CircuitPython for IoT Devices

CircuitPython and Raspberry Pi Pico can together transform everyday devices into smart wonders. This short tutorial will help you write a simple program for using CircuitPython and upload it to Raspberry Pi Pico.

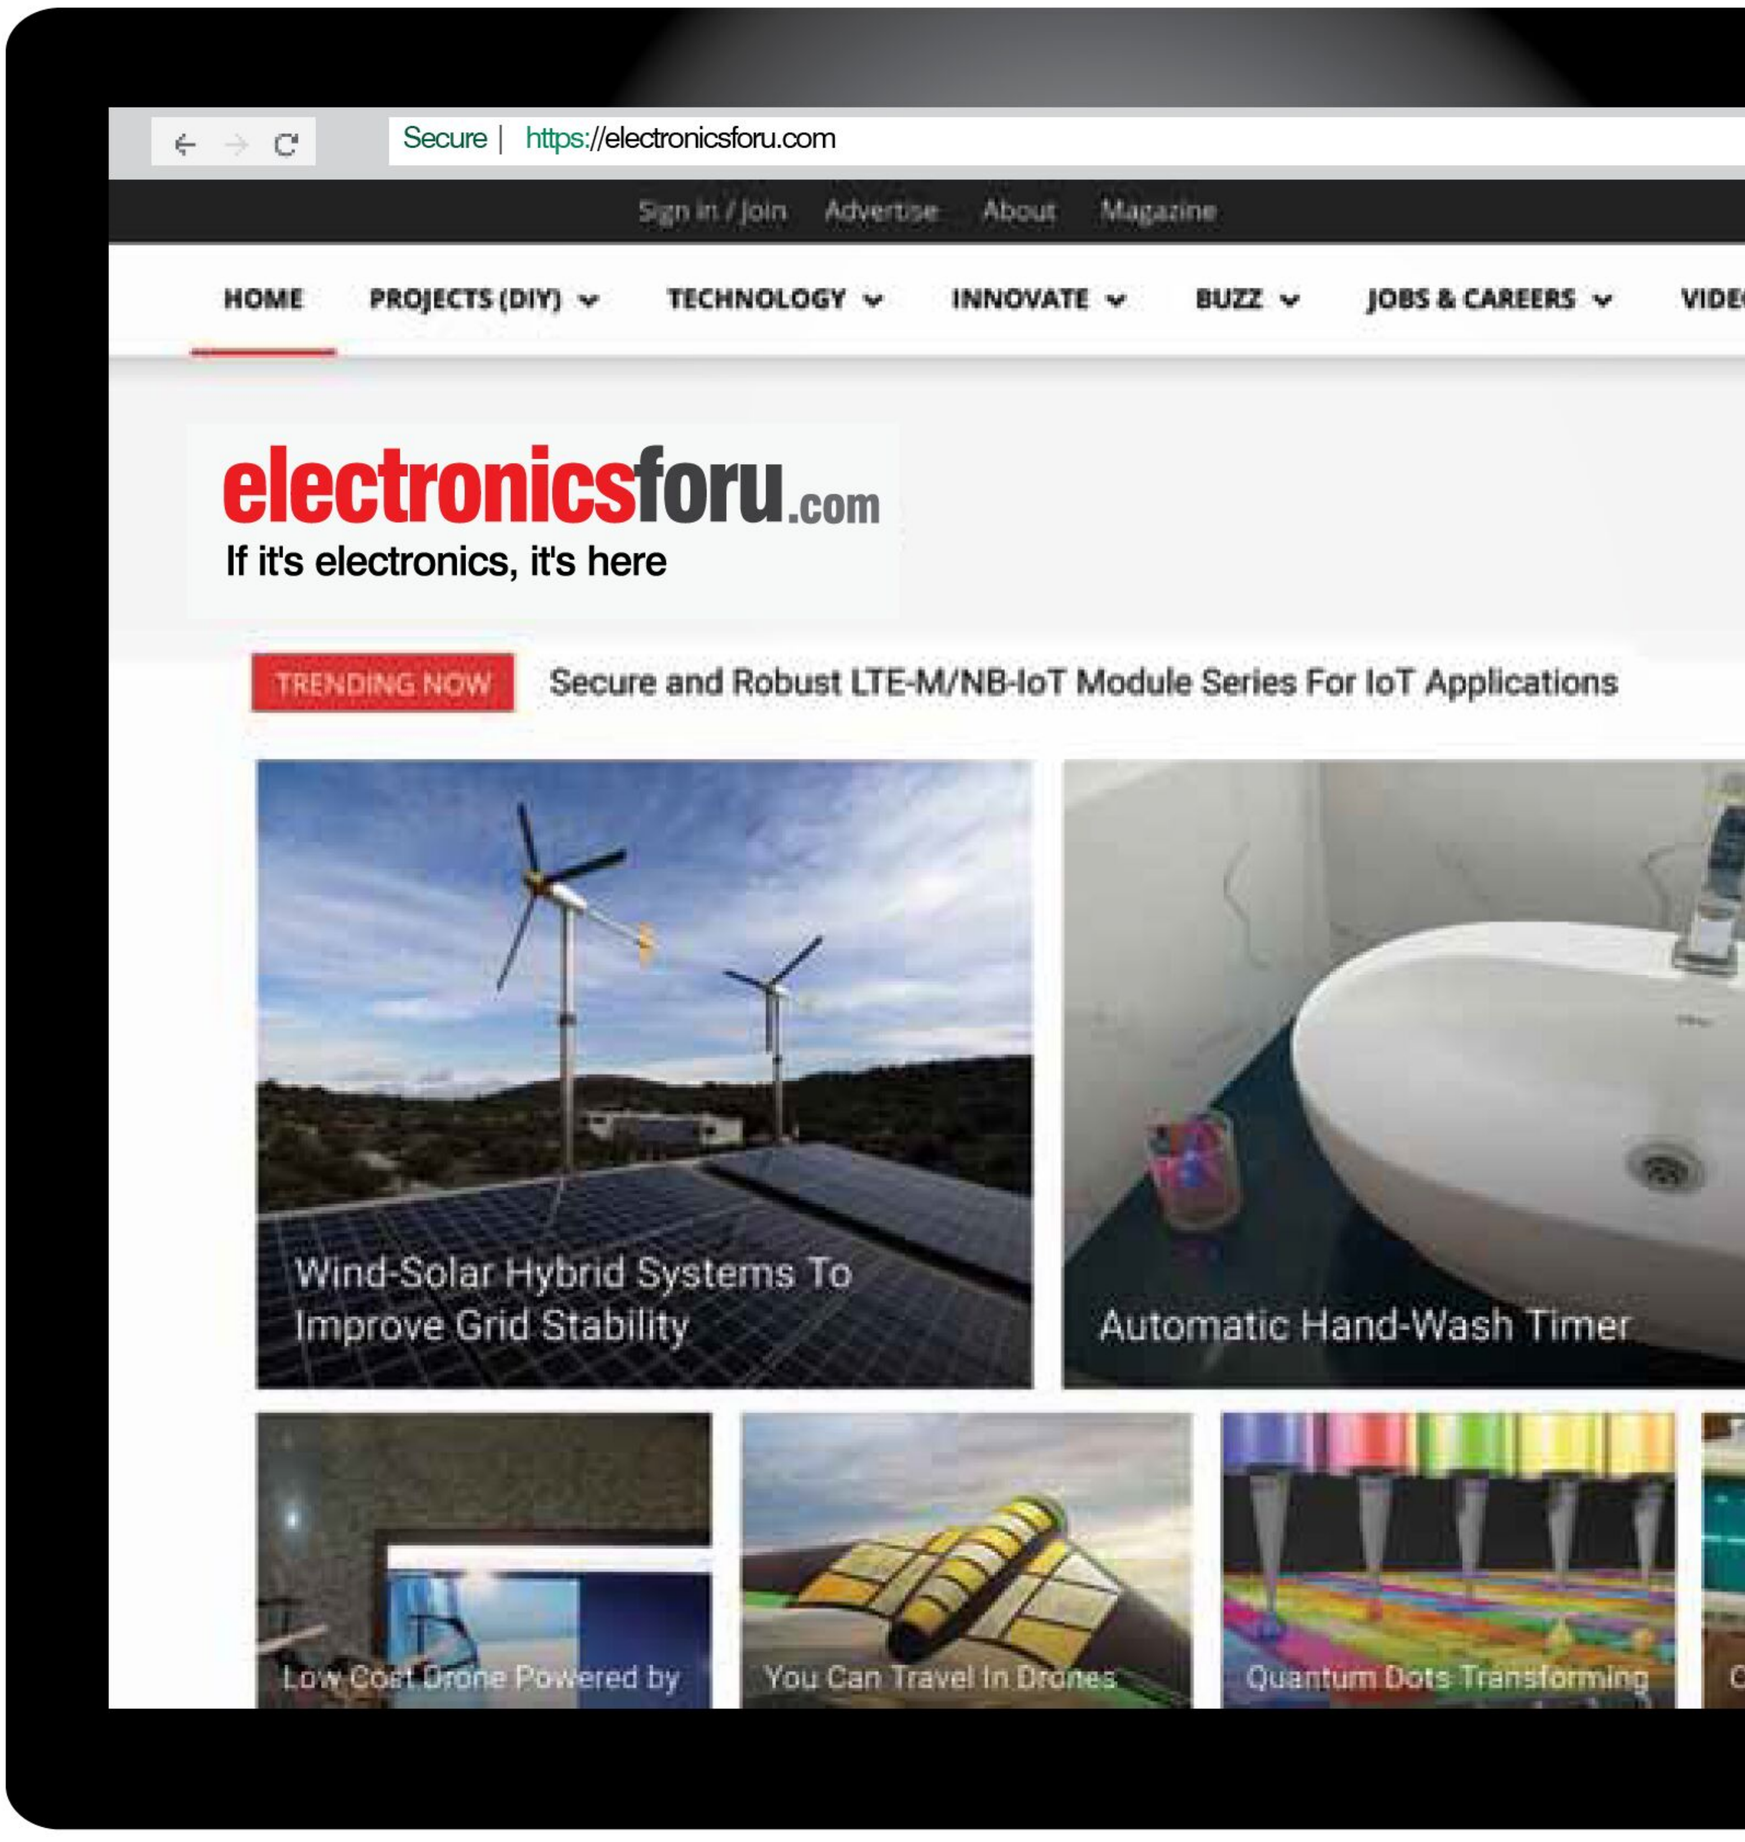
Read more on page.....73

Choose the Right Cloud Platform for Implementing IoT Protocols

Working with IoT protocols like MQTT, AMQP, and CoAP on cloud platforms is essential for developing scalable and efficient IoT applications. The choice of the programming platform will depend on factors like project requirements, existing skills, and target devices. Leveraging the appropriate libraries and cloud services can enable seamless integration of IoT devices with cloud-based applications.

Read more on page.....75

Your favourite website has



electronicsforu.com

THANKS TO YOU—OUR ONLINE NETWORK IS

FACTS & FIGURES

- 4 websites (two more coming soon)
- Five major Facebook communities
- Seven major LinkedIN groups & pages
- Million-plus active users (monthly)
- Million-plus reach through Facebook
- Fifty-thousand-plus industry connections through LinkedIN

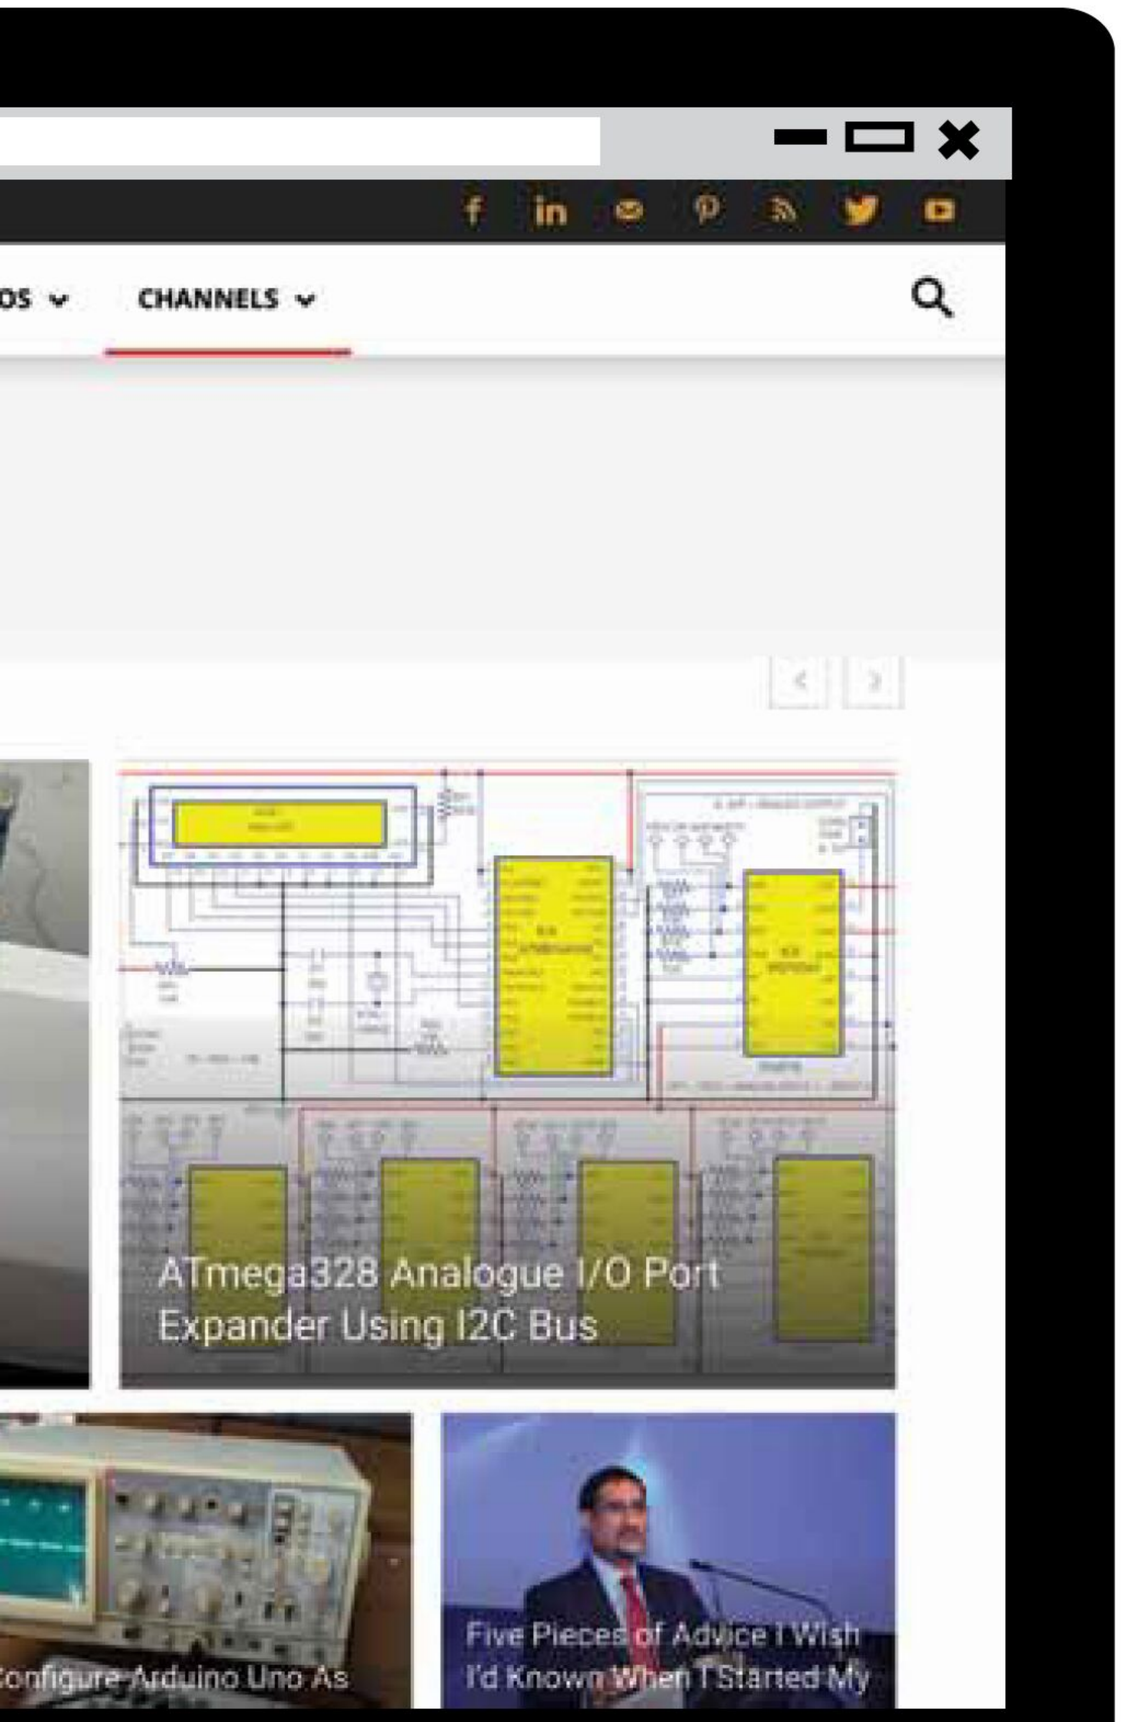
READERS

- You can access all content for FREE
- You can subscribe to newsletters for FREE --on most websites
- Register on our websites to get free invites to technical webinars and seminars

EXPERTS

- Experts who want to share their knowledge through articles, DIY Projects, etc are welcome
- We also welcome experts who want to share their knowledge through webinars or seminars
- You can contact us at editop@efy.in

fast growing peers now...



Amazing DIY Projects. Latest Tech trends.

The hang-out for electronics enthusiasts.



The Latest in IOT.
A platform for enablers, creators and providers of IOT solutions.



India. Electronics. Directory.
Enabling commerce between buyers & sellers of electronics in India.



Business. Electronics. India.
Everything you want to know about India's electronics industry.

AMONGST THE WORLD'S TOP 5 AND GROWING!

INDUSTRY

- You can advertise for as little as US$ 100 per month
- Special combo offers for advertisers in our print publications
- We've now enabled flexible CPM-based advertising
- You can advertise on the platform of your choice (based on your target audience)
- We invite press releases at efy-edit-team@efy.in
- Press releases are published free of cost, subject to discretion of the editorial team

RESPONSE GUARANTEED SOLUTIONS

We now also act as marketing partners for our clients and drive entire marketing for them, where we charge them on basis of results and not efforts!

CONTACT US: Shrikant Rao • growmybiz@efy.in • +91-98111 55335

Generative AI in IoT: Moving Towards Transformative Applications

Generative AI's contribution to the Internet of Things (IoT) cannot be underestimated. It has the potential to truly transform this domain by improving the security of IoT networks and making IoT devices more affordable, among other things.



Since the launch of ChatGPT 3.5 on November 30, 2022, no one seems to be able to stop talking or hearing about the buzzwords 'generative AI'. Generative AI is a type of artificial intelligence (AI) that can generate new content, such as text, images, audio, and video. It is trained on large data sets of existing content, and then uses this knowledge to create new and original content. Just like everything else, generative AI has the potential to transform the Internet of Things (IoT) in several ways. For example, it can be used to create synthetic data for training machine learning models, generate personalised experiences for users, or improve anomaly detection, and much more.

Let's explore some of the specific ways in which generative AI is being used in IoT today, as well as review some of the potential future applications of this technology.

Creating synthetic data for machine learning

Leveraging generative AI to craft synthetic data for training machine learning models tackles a significant challenge in developing such models for IoT devices. While IoT devices can generate a lot of data, collecting and labelling it often proves expensive and time-consuming. Synthetic data originates artificially rather than from real world sources. It can be used to

train machine learning models for a variety of tasks, such as anomaly detection, predictive maintenance, and quality control. For example, a manufacturing company could use generative AI to create synthetic data of machine failure scenarios. This data could then be used to train a machine learning model to detect anomalies in real-time and predict when machines are likely to fail.

Crafting personalised experiences

Generative AI can deliver personalised experiences to users of IoT devices. A smart home system can use generative AI to curate personalised lighting and temperature settings for individual users. Likewise, wearable devices can use generative AI to offer personalised workout recommendations. To create these personalised experiences, generative AI algorithms need to be trained on data about the user's preferences. This data can be collected from a variety of sources, such as surveys, user interactions with the IoT device, and data from other sensors.

Improving anomaly detection

Improving anomaly detection with generative AI entails the identification of unusual or unexpected events in data. It is an important task for IoT devices, as it can be used to identify problems such as machine failures, security breaches, and fraud. Generative AI can be used to improve anomaly detection by generating synthetic data that represents normal operating conditions. This data can then be used to train a machine learning model to identify anomalies in real-time. For example, a power grid operator could use generative AI to create synthetic data of normal power consumption patterns. This data could then be used to train a machine learning model to detect anomalies in real-time, such as sudden spikes in power consumption.

Enabling on-device machine learning

On-device machine learning empowers IoT devices to run machine learning models locally without relying on cloud resources. Generative AI can be used to enable on-device machine learning by generating smaller and more efficient machine learning models. This is important because IoT devices often have limited computing resources. For example, a company could use generative AI to generate a smaller and more efficient machine learning model for anomaly detection. This model could then be deployed on IoT devices to enable on-device anomaly detection.

Automating network management

The complexity of automating network management for IoT networks requires innovative solutions. Generative AI emerges as a valuable solution to automate various aspects of network management such as device configurations and network traffic optimisation. For example, a company could use generative AI to create a system that automatically configures new devices when they are added to the network. This system could also be used to optimise network traffic by routing data over the most efficient paths.

Potential future IoT applications

As generative AI is a rapidly developing field, there are many potential future applications for this technology in IoT. For example, generative AI could be used to:

- Create new types of IoT devices, such as smart assistants that can understand and respond to natural language
- Develop new ways to interact with IoT devices, such as using gestures or voice commands
- Improve the security and reliability of IoT networks
- Make IoT devices more affordable and accessible to everyone

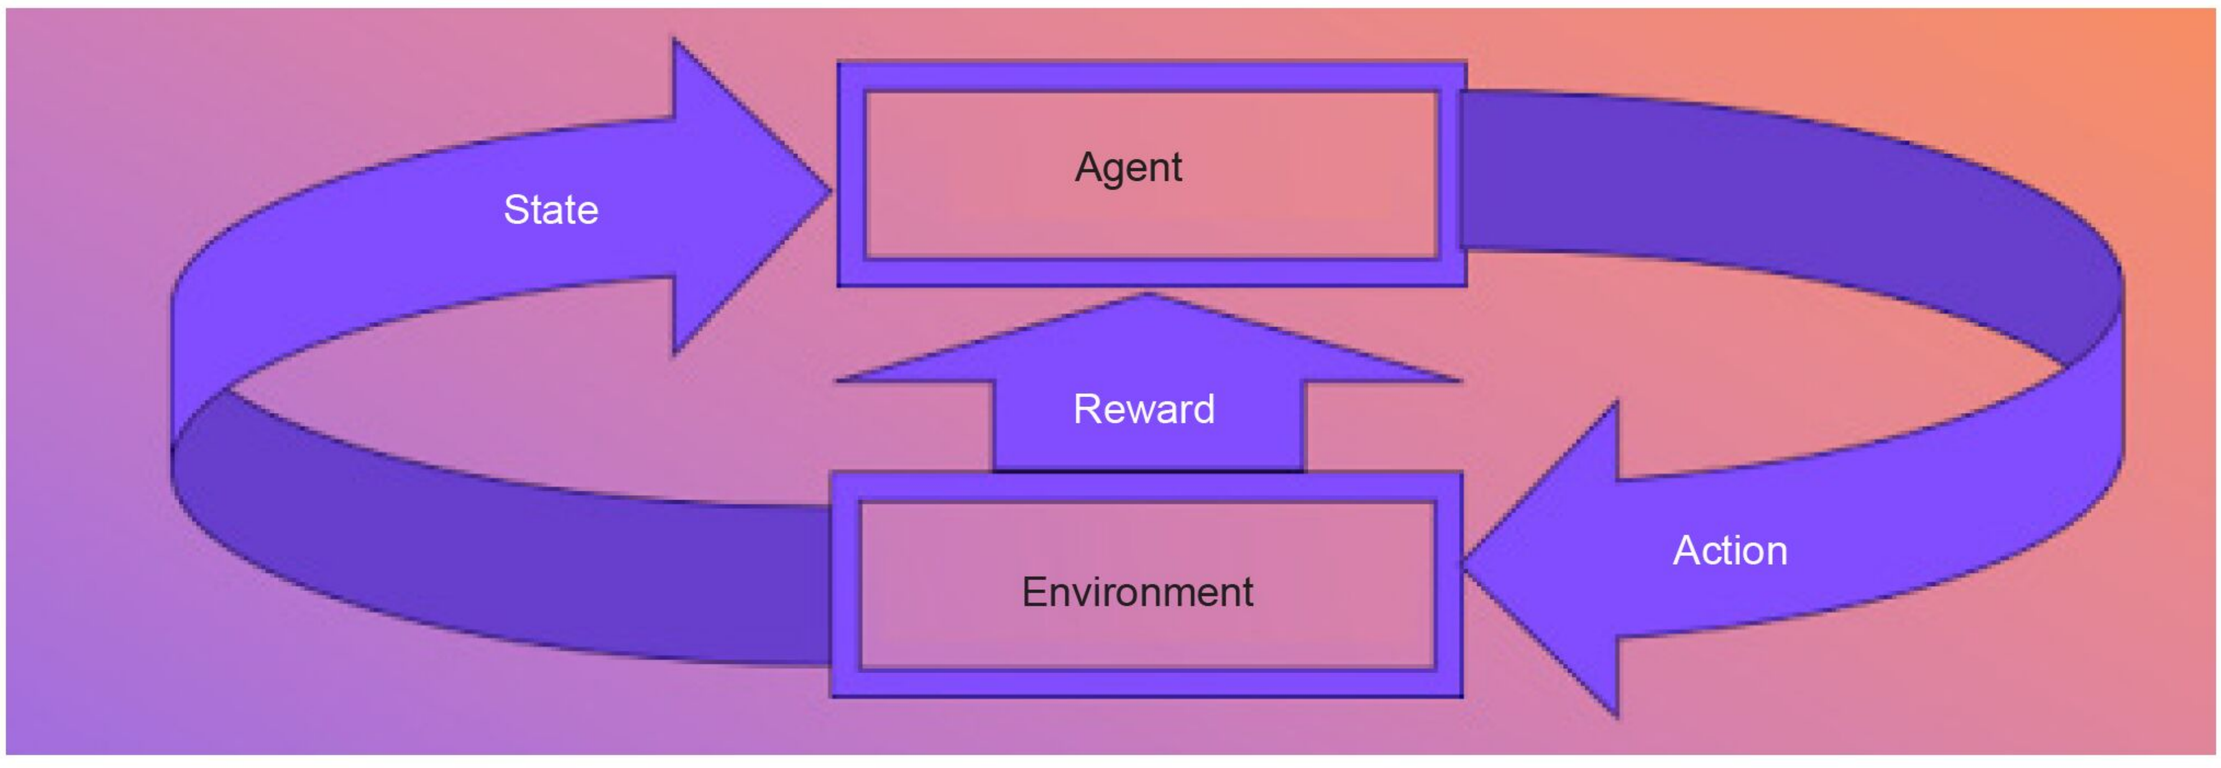


Figure 1: OpenAI Gym: reinforcement learning

OpenAI Gym

As generative AI becomes mainstream, the open source community is contributing significantly to the journey. OpenAI is the front runner in advancing the innovation, education, and adoption of the newfound capabilities. OpenAI Gym is a toolkit for developing and comparing reinforcement learning algorithms. It provides a set of environments in which agents can learn to perform tasks such as balancing a pole, walking, and playing games.

OpenAI Gym is based on the Markov Decision Process (MDP) framework, which is a mathematical model for decision-making in sequential environments. In an MDP, an agent interacts with an environment by taking actions and receiving rewards. The agent's goal is to learn to take actions that maximise its expected reward over time.

OpenAI Gym environments are typically simulated, but they can also be connected to real-world hardware. This allows agents to be trained in a variety of settings and then deployed to the real world without having to retrain them.

Let's look at OpenAI Gym environments and how they work.

- ***State:*** The state of the environment represents the agent's current situation. For example, the state of the CartPole environment includes the position and velocity of the cart, and the angle and angular velocity of the pole.
- ***Action space:*** The action space is the set of actions that the agent can take. For example, the action space in the CartPole environment consists of two actions: pushing the cart left or right.
- ***Reward function:*** The reward function specifies the reward that the agent receives for taking a particular action in each state. For example, the agent receives a reward of +1 for keeping the pole balanced in the CartPole environment.

To learn how to perform a task in an OpenAI Gym environment, an agent uses a reinforcement learning algorithm. Reinforcement learning algorithms work by trial and error. The agent tries different actions and observes the rewards that it receives. The agent then uses this information to update its policy, which is a mapping from states to actions.

Over time, the agent learns to take actions that maximise its expected reward.

OpenAI Gym provides a variety of environments for training reinforcement learning agents. Some of the most popular environments include:

- *CartPole:* The agent must balance a pole on a moving cart
- *Acrobot*: The agent must control a two-jointed robot arm
- *LunarLander:* The agent must land a lunar lander on a platform
- *BipedalWalker:* The agent must walk a bipedal robot
- *Roboschool:* A suite of environments for training robot locomotion skills

OpenAI Gym has a number of benefits.

- ***Easy to use:*** OpenAI Gym is easy to use, even for beginners. It provides a simple API for interacting with environments and training agents.
- ***Flexible:*** It is flexible and can be used to train agents for a variety of tasks. It also supports a variety of reinforcement learning algorithms.
- ***Community:*** OpenAI Gym has a large and active community. This means that there are a lot of resources available to help users get started and troubleshoot problems.

Generative AI has the potential to revolutionise the Internet of Things (IoT) in a number of ways. It is still a developing field, but can transform the way we interact with the world around us. As generative AI technology continues to mature, we can expect to see even more innovative and transformative applications in IoT and beyond. Open source solutions such as OpenAI Gym can be used to develop and train generative AI models for IoT applications. This makes generative AI more accessible to businesses and individuals of all sizes. The future of generative AI in IoT is bright. As technology continues to develop and mature, we can expect to see even more innovative and transformative applications. END



References

- OpenAI Gym project in GitHub; *https://github.com/openai/gym*
- Gymnasium: An API standard for reinforcement learning with a diverse collectionof reference environment; *https://gymnasium.farama.org/*
- Using generative AI for IoT; *https://iot-analytics.com/3-generative-aiot-applications-beyond-chatgpt/*
- Generative Artificial Intelligence at the Edge in the Modern Internet of Things; *https://www.comsoc.org/publications/magazines/ieee-internet-things-magazine/cfp/generative-artificial-intelligence-edge*

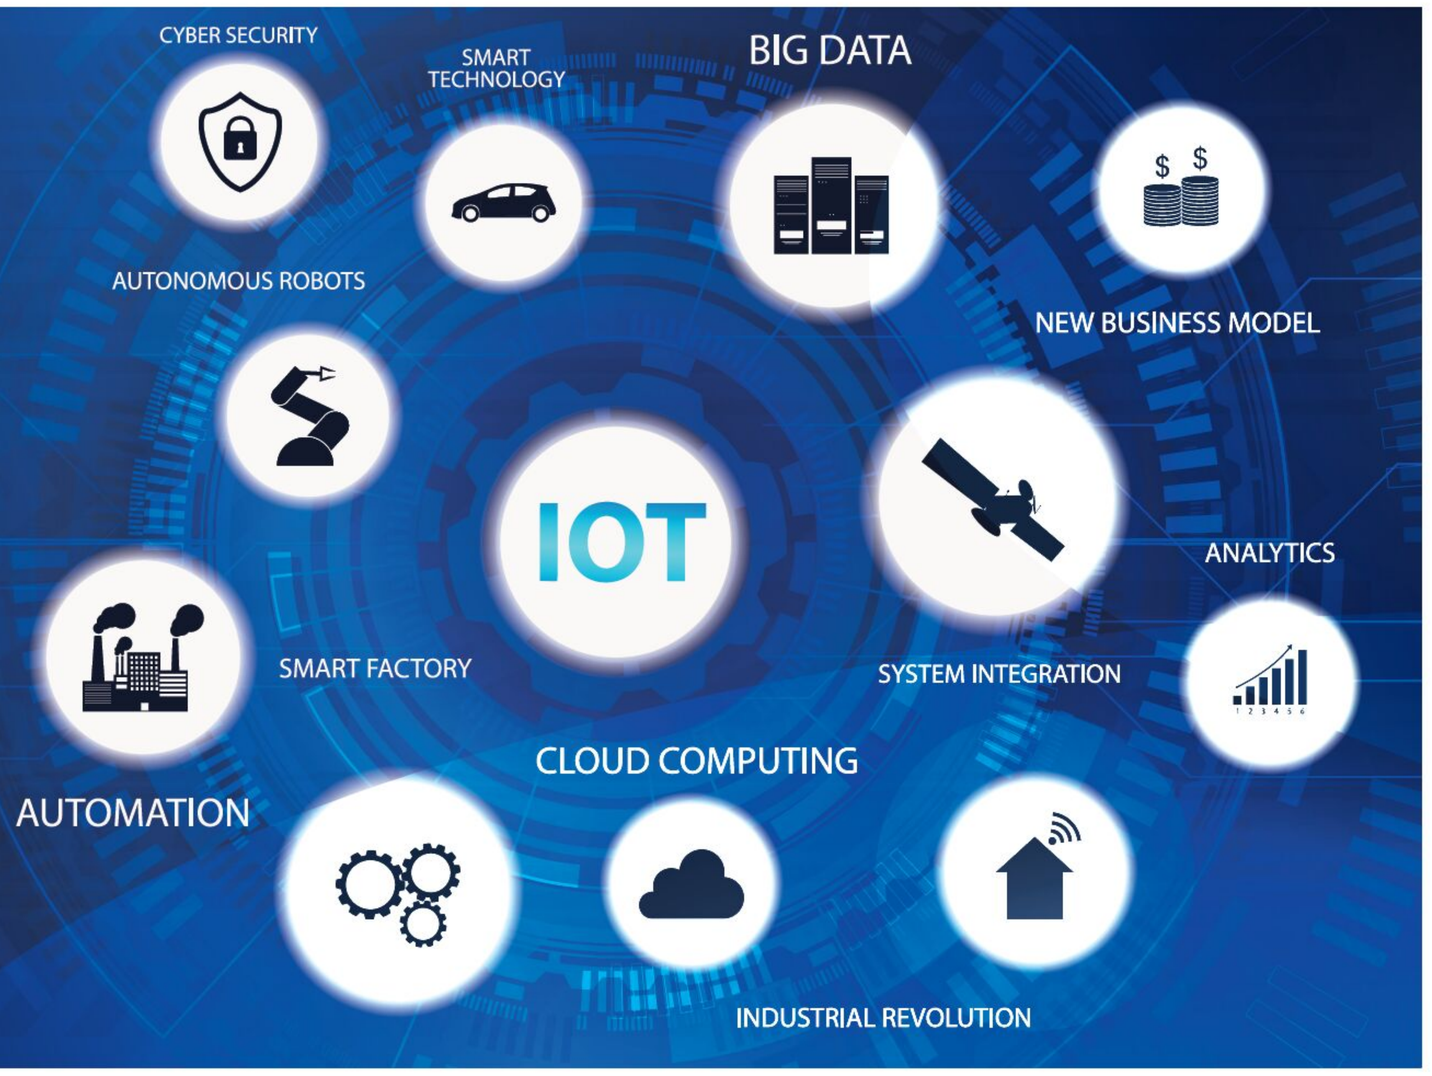
By: Bala Kalavala

The author is a technical architect, evangelist, thought leader, and sought-after keynote speaker. He currently works as digital transformation business development executive in a global technology consulting firm.

Disclaimer: This article expresses the views of the author and not of the organisation he works in.

The Role of Open Source in the Growth and Adoption of IoT

The symbiotic relationship between open source and IoT is undeniable. Open source technologies have been instrumental in fuelling IoT's growth and widespread adoption. This relationship is underpinned by the principles of accessibility, flexibility, interoperability, security, and innovation that open source brings to the IoT landscape.



The world is currently experiencing a profound transformation driven by the Internet of Things (IoT), which allows physical objects, devices, and sensors to gather, share, and act upon data via the internet. In this article, we delve into the importance of open source software in fostering the development and widespread acceptance of IoT solutions.

At its core, IoT represents the interconnectivity of everyday objects through the internet. This connectivity facilitates real-time data collection, analysis, and utilisation, revolutionising industries ranging from healthcare to agriculture, and manufacturing to smart cities.

In agriculture, IoT sensors monitor soil conditions and crop health, optimising resource utilisation. In healthcare, wearable devices enable remote patient monitoring, leading to improved health outcomes. Manufacturing benefits from predictive maintenance powered by IoT, reducing downtime and costs. These examples underscore the pervasive impact of IoT, spanning diverse sectors.

Challenges faced in IoT implementation

Despite its potential, the implementation of IoT projects presents a set of unique challenges, including:

- ***Security:*** Protecting IoT devices from cyberattacks
- ***Interoperability:*** Ensuring devices can work together
- ***Scalability:*** Handling a growing number of devices
- ***Data management:*** Efficiently handling large data volumes
- ***Privacy:*** Managing personal data and privacy concerns
- ***Power consumption:*** Optimising energy usage for battery-powered devices
- ***Reliability:*** Ensuring devices work in various conditions
- ***Costs:*** Dealing with initial and ongoing expenses
- ***Skills gap:*** Finding and retaining IoT talent
- ***Regulatory compliance:*** Meeting industry-specific regulations
- ***Integration:*** Merging IoT with existing systems
- ***Environmental impact:*** Sustainable device disposal
- ***Ethical considerations:*** Addressing data ethics and usage

The need for cost-effective and flexible solutions

Cost-effective and flexible solutions are indispensable in the IoT landscape. They ensure widespread adoption, scalability, efficient resource use, and competitiveness. These solutions accommodate innovation, support global reach, empower startups, lower

total ownership costs, and promote sustainability, fostering IoT's growth across diverse industries and applications.

While initial investment requirements may exist, cost-effective IoT solutions help minimise ongoing maintenance and operational expenses. Consequently, these solutions offer long-term benefits by reducing the total cost of ownership (TCO). Moreover, as sustainability gains increasing importance, IoT solutions contribute by reducing energy consumption and electronic waste, aligning with broader environmental goals.

Open source and IoT: A perfect match

Open source software empowers IoT developers to leverage the power of community-driven innovation while avoiding vendor lock-in.

Here are some advantages of using open source in IoT projects.

- *Flexibility and customisation:* IoT solutions may be easily modified to suit the specific project requirements.
- *Cost-effectiveness:* Open source eliminates or reduces expensive licensing costs.
- *Community support:* The open source community provides vast knowledge and resources for developers to seek assistance, share insights, and collaborate on IoT projects, fostering innovation and problem-solving.

Real-world examples of successful open source IoT projects

Several open source IoT projects have made a significant impact in the real world. Here are some notable examples.

- ***Eclipse IoT:*** The Eclipse IoT project is a community-driven effort providing open source IoT technologies and frameworks. It includes projects like Eclipse Paho (MQTT client), Eclipse Mosquitto (MQTT broker), and Eclipse Kura (IoT gateway). These tools are widely used in building scalable and interoperable IoT solutions.
- ***OpenHAB***: OpenHAB is a home automation platform that can be used to manage and coordinate a wide range of smart home technologies. It supports a wide range of IoT protocols and devices, making it a popular choice for DIY smart home enthusiasts.
- ***Home Assistant:*** Home Assistant is another open source home automation platform known for its user-friendly interface and extensive support for IoT technology. It allows users to create custom automation rules and integrates with numerous smart home devices and services.
- ***Node-RED:*** Node-RED is a flow-based development tool for the visual programming of IoT applications. It simplifies the creation of IoT workflows by providing a drag-and-drop interface for connecting devices, APIs, and online services.
- ***Zephyr Project:*** The Zephyr Project is an open source, real-time operating system (RTOS) designed for resource-constrained IoT devices. It offers a lightweight and scalable platform for developing IoT applications in various industries.
- ***Tasmota:*** Tasmota is the firmware for ESP8266 and ESP32-based IoT devices, allowing users to flash their devices with open source software. It provides greater control and compatibility for devices like smart switches, lights, and sensors.
- ***ThingsBoard:*** ThingsBoard is an open source IoT platform that offers data visualisation, device management, and rule engine capabilities. It enables users to collect, process, and display IoT data in real time.
- ***RIOT:*** RIOT is an open source operating system designed for IoT. It is particularly well-suited for low-power, resource-constrained IoT devices.
- ***Cayenne by myDevices:*** This is an intuitive IoT platform with drag-and-drop widgets for sensor data visualisation and device control.
- ***Grafana:*** Grafana is a popular open source analytics and monitoring platform often used for visualising data from IoT devices and sensors.

Pros and cons of using open source hardware for IoT development

Pros

- ***Affordability:*** Open source hardware is often cost-effective, making it accessible for hobbyists, startups, and educational institutions.
- ***Customisability:*** Open source hardware can be modified and adapted to specific project requirements. This flexibility enables developers to create tailored IoT devices for unique applications.
- ***Community support:*** Many open source hardware platforms have active communities of developers and enthusiasts. This support network provides valuable resources, including documentation, tutorials, and forums for troubleshooting.
- ***Transparency:*** Open source hardware designs are typically open and transparent. Developers can access detailed schematics and design files, allowing for in-depth understanding and modification of the hardware.
- ***Interoperability:*** Open source hardware often adheres to open standards, enhancing interoperability with other devices and software.
- ***Rapid prototyping:*** Open source hardware platforms facilitate rapid prototyping, enabling developers to quickly test and iterate on IoT ideas and concepts.

Cons

- ***Resource constraints:*** Some open source hardware platforms may have limitations in terms of processing power, memory, or connectivity options.

- ***Hardware availability:*** Availability of specific open source hardware components may vary by region or over time.
- ***Support and longevity:*** While many open source hardware platforms have active communities, the longevity of community support can vary.
- ***Integration challenges:*** Integrating open source hardware with other proprietary systems or software may present challenges.
- ***Complexity:*** Some open source hardware platforms may have a steeper learning curve, particularly for beginners, requiring adequate training and resources to get started.
- ***Certifications:*** For IoT applications requiring regulatory compliance (e.g., safety or wireless certifications), open source hardware may require additional validation and testing.

IoT communication protocols and their open source alternative

IoT communication protocols are crucial for efficient and secure data exchange. Some common IoT communication protocols and their open source alternatives include:

MQTT (Message Queuing Telemetry Transport): A lightweight and efficient publish-subscribe messaging protocol widely used in IoT.

Open source alternative: Eclipse Paho provides open source MQTT client libraries for various programming languages, facilitating MQTT communication in IoT applications.

CoAP (Constrained Application Protocol): This protocol is designed for resource-constrained IoT devices and is suitable for low-power, constrained environments.

Open source alternative: Californium is an open source implementation of CoAP for Java-based IoT applications, enabling efficient communication between devices.

HTTP/HTTPS (Hypertext Transfer Protocol): While not exclusive to IoT, HTTP and HTTPS are commonly used for IoT device communication.

Open source alternative: Web servers like Apache HTTP Server and NGINX are open source and often utilised for handling HTTP(S) traffic in IoT applications.

WebSockets: This protocol enables real-time, bi-directional communication between clients and servers.

Open source alternative: Libraries like Socket.io provide open source WebSocket support for IoT applications, facilitating real-time data exchange.

LoRaWAN (Long Range Wide Area Network): This is a wireless communication protocol for long-range, low-power, IoT devices.

Open source alternative: The Things Network (TTN) is an open source, community-driven network supporting LoRaWAN deployments. It provides open access to LoRaWAN connectivity.

Bluetooth Low Energy (BLE): BLE is commonly used for short-range communication between IoT devices and smartphones.

Open source alternative: The Zephyr Project is an open source real-time operating system (RTOS) that supports BLE communication, enabling the development of custom BLE-based IoT solutions.

XMPP (Extensible Messaging and Presence Protocol): This open communication protocol is often used in IoT for real-time messaging and presence information.

Open source alternative: ejabberd is an open source XMPP server that facilitates secure, scalable, and extensible communication for IoT applications.

DDS (Data Distribution Service): This is a protocol for real-time data exchange between IoT devices, suitable for applications requiring high reliability.

Open source alternative: Eclipse Cyclone DDS is an open source implementation of the DDS protocol, providing a framework for building robust and efficient IoT systems.

Benefits of open source in IoT

- ***Cost-effectiveness and reduced licensing fees***: Open source IoT solutions are typically cost-effective as they eliminate or significantly reduce licensing fees. This minimises the financial barriers to IoT development and deployment, allowing for wider adoption and innovation.
- ***Flexibility and customisation for specific IoT applications:*** Open source IoT software and hardware offer a high degree of flexibility and customisation. This adaptability ensures that IoT technology can be finely tuned to address unique use cases across healthcare, agriculture, industrial automation, and more industries.
- ***Collaboration and community support in open source IoT projects:*** Open source fosters collaboration among developers, researchers, and enthusiasts. IoT projects that embrace open source principles benefit from a collective effort to improve and enhance the technology. Communities formed around open source IoT projects offer valuable resources, including documentation, tutorials, forums, and peer-to-peer support. This collaborative ecosystem accelerates IoT development, encourages knowledge sharing, and helps address challenges effectively.

Case studies

Let's take a quick look at a few real-world IoT applications based on open source.

Smart agriculture with open source IoT

- **Project overview:** A large-scale smart agriculture project aimed to optimise crop production and resource management on a farm.
- **Open source solution:** The

project leveraged open source IoT platforms like Eclipse IoT and open source hardware like Raspberry Pi and Arduino for data collection and management.

Challenges faced:

- **Interoperability:** The farm had a mix of sensors and devices from different manufacturers. Ensuring interoperability and data integration was a challenge.
- **Cost management:** The project had budget constraints, making it essential to find cost-effective solutions.

Success factors:

- **Eclipse IoT:** Eclipse IoT provided a unified platform for connecting diverse devices and sensors, addressing the interoperability challenge.
- **Community support:** The open source community provided guidance and solutions, helping the team overcome technical hurdles.
- **Customisation:** Open source allows for tailored solutions, ensuring efficient resource management based on specific farm requirements.

Industrial IoT in manufacturing

- **Project overview:** A manufacturing plant sought to implement Industrial IoT for predictive maintenance and process optimisation.
- **Open source solution:** The project used open source IoT platforms like Node-RED and open source hardware for data acquisition.

Challenges faced:

- **Legacy equipment:** Integrating IoT with legacy machinery and systems presented compatibility challenges.
- **Data security:** Ensuring the security of sensitive manufacturing data was a top concern.

Success factors:

- **Node-RED:** Node-RED's visual programming capabilities simplified the integration of IoT with existing systems.
- **Community expertise:** The open source community offered security best practices and solutions to safeguard data.
- **Cost savings:** Open source solutions minimised licensing costs, allowing the project to allocate resources effectively.

Smart city IoT deployment

- **Project overview:** A city government aimed to create a smart city infrastructure for improved public services and urban planning.
- **Open source solution:** The project adopted open source IoT platforms like Eclipse IoT and Combin, and technology for long-range connectivity.

Challenges faced:

- **Scalability:** Deploying IoT sensors and gateways across the entire city required a scalable and cost-effective solution.
- **Data management:** Managing and analysing vast amounts of sensor data for urban planning posed a significant challenge.

Success factors:

- **LoRaWAN:** LoRaWAN's long-range capability reduced the number of gateways needed, addressing scalability concerns.
- **Eclipse IoT:** Eclipse IoT's data management and analytics tools enabled efficient data processing and urban insights.
- **Open standards:** The use of open standards ensured interoperability and flexibility in expanding the smart city infrastructure.

Addressing the challenges

Addressing challenges when using open source IoT is crucial to ensure the success and security of the projects. Here are some key points to consider.

Potential challenges of using open source in IoT and how to address them

- **Interoperability:** The IoT ecosystem often comprises diverse devices and protocols. Ensuring interoperability between open source components and proprietary systems can be challenging. Implementing open standards and protocols can help mitigate this challenge.
- **Community support:** While open source communities can be a valuable resource, the level of support may vary. Organisations should assess the strength and responsiveness of the community behind the chosen open source projects.
- **Integration complexity:** It might be difficult to combine open source IoT solutions with pre-existing IT frameworks and legacy systems. A well-defined integration strategy and architecture are essential to overcome this challenge.

Security concerns and how to mitigate them

- **Vulnerabilities:** Open source software may have vulnerabilities that can be exploited. Regularly monitoring for security updates and applying patches promptly is critical.
- **Authentication and access control:** IoT infrastructure and data can be protected from hackers by using multi-factor authentication and other access restrictions. Open source solutions often provide security features but require proper configuration.
- **Data encryption:** Data transmitted between IoT devices and servers must be encrypted using secure protocols like TLS/SSL to protect against eavesdropping and data breaches.
- **Device security:** IoT devices should be secured by regularly updating their firmware, changing default passwords, and implementing

device-level security measures.

Scalability and maintenance considerations

- **Scalability planning:** IoT solutions must be designed to handle growth. Consider how open source components can scale with increasing device deployments and data volume. Evaluate cloud-based solutions if needed.
- **Data management:** As IoT generates vast amounts of data, consider how data will be stored, processed, and analysed. Scalable databases and data analytics tools may be necessary.
- **Maintenance:** Regularly maintain and update open source components to ensure they remain secure and compatible with evolving standards. Establish a clear maintenance schedule and procedures.
- **Vendor lock-in:** Avoid vendor lock-in by selecting open source technologies with a strong community and a clear roadmap. This ensures long-term support and prevents reliance on a single vendor.

The future of IoT and open source

Predictions for the future of IoT adoption

- **Massive expansion:** IoT adoption is expected to grow exponentially across industries, including healthcare, agriculture, smart cities, and manufacturing. The proliferation of connected devices will contribute to the digital transformation of businesses and societies.
- **Edge computing dominance:** Edge computing will play a pivotal role in IoT as organisations seek to process data closer to the source. This will reduce latency, improve real-time decision-making, and enhance security.
- **AI and machine learning integration:** IoT devices will increasingly incorporate AI and machine learning capabilities, enabling them to analyse data locally and make intelligent decisions autonomously. This trend will lead to more efficient IoT systems.

How open source is likely to evolve in the IoT space

- **Greater proliferation:** Open source will continue to play a significant role in IoT development due to its accessibility, flexibility, and cost-effectiveness. More open source IoT projects and frameworks will emerge, covering a broader range of applications.
- **Standardisation and interoperability:** Open source initiatives will focus on standardisation and interoperability, aiming to create common protocols and frameworks that facilitate seamless communication between IoT devices and platforms.
- **Security enhancements:** Open source security tools and practices will evolve to address the growing threat landscape in IoT. Enhanced security measures will become standard, including secure boot mechanisms and device identity management.
- **Edge and fog computing:** Open source projects for edge and fog computing will become more prevalent, empowering developers to build efficient, decentralised IoT solutions that reduce the load on centralised cloud servers.

Opportunities for further growth and innovation

- **Industry-specific solutions:** Open source IoT will enable the creation of industry-specific solutions, such as precision agriculture, personalised healthcare, and smart manufacturing. These tailored solutions will drive innovation and efficiency within sectors.
- **Open hardware ecosystem:** The growth of open source hardware, like Raspberry Pi and Arduino, will continue. This will provide hardware developers with a strong foundation for building IoT devices, reducing development costs, and fostering innovation.
- **Data monetisation:** Organisations will explore new ways to monetise IoT-generated data, including data marketplaces and data-as-a-service models. Open source platforms will facilitate data sharing while preserving privacy and security.
- **Cross-domain integration:** IoT and open source will converge with other technologies like blockchain and AI to create integrated, cross-domain solutions. For example, blockchain can enhance IoT device identity and data integrity.

Open source is the cornerstone of IoT development, offering the tools and frameworks necessary to create scalable, secure and customised solutions. By embracing open source, businesses and developers can overcome barriers, reduce costs, enhance security, and tap into a vibrant community of innovators. Open source is not merely an option; it is the pathway to unlocking the transformative potential of IoT, driving progress, and creating a more interconnected and intelligent world. END

By: Sufyan bin Uzayr

The author, apart from being a writer, is an educator and entrepreneur. He holds multiple degrees in management, information technology and humanities. He has written seven books and is currently the director of a multinational IT company. As an educator, he has taught at universities and institutions worldwide.

How Open Source Tools Help with Real-time Data Analytics in IoT

Real-time data analytics in IoT offers immense potential to optimise processes, enhance user experiences, and drive meaningful insights. Open source tools play a pivotal role in realising this potential, providing cost-effective, flexible, and powerful solutions for real-time analytics.

The rapid evolution of technology has ushered us into the era of the Internet of Things (IoT), where an extensive network of interconnected devices is generating an unprecedented volume of data. This data holds immense potential, offering insights into consumer behaviour, operational efficiency, and numerous other aspects that can drive informed decision-making. However, the true value of this data lies in its timely analysis.

Real-time data analytics involves the immediate processing and analysis of data as it is generated. This instantaneous approach allows for proactive responses, identifying patterns, trends, and anomalies in real time. When applied to IoT, real-time analytics enables organisations to optimise operations, enhance user experiences, and even prevent critical failures.



Understanding real-time data analytics in IoT

IoT is not merely about connecting devices; it's about leveraging the data these devices produce. Consider a scenario where an urban traffic management system utilises IoT sensors embedded in traffic lights, cameras, and vehicles to collect real-time traffic data. Without real-time analytics, this data remains a vast, untapped resource. However, with real-time analytics in place, the system can instantly analyse traffic patterns, detect congested areas, and dynamically adjust traffic light timings, optimising traffic flow.

The volume, velocity, and variety of data generated by IoT devices necessitate real-time analytics. Traditional data processing methods fall short in handling this immense influx of data in a timely manner. Real-time analytics, on the other hand, empowers organisations to make decisions within milliseconds or seconds of data generation.

Open source tools for real-time data analytics

To harness the power of real-time analytics in IoT, open source tools play a pivotal role. These tools are not only cost-effective but also offer a high degree of flexibility and customisation to suit specific IoT requirements.

Apache Kafka: Managing data streams

Apache Kafka is a leading open source platform designed for handling real-time data streaming. It acts as a high-throughput, fault-tolerant, and horizontally scalable platform that allows you to publish and subscribe to streams of records. Think of it as a messaging system on steroids, well-suited for real-time analytics.

Example: How Uber uses Kafka to process and analyse ride data in real-time

Uber, the ride-sharing giant, employs Apache Kafka to process an enormous amount of data generated by users, drivers, and vehicles in real time. Every trip, GPS coordinate, and user interaction is captured and streamed through Kafka, enabling Uber to optimise ride matching, detect fraud, and enhance overall service efficiency.

Apache Flink: Stream processing

Apache Flink is an open source stream processing framework that provides robust support for event-driven applications. It excels in handling high-throughput, low-latency streaming analytics. Flink processes data in real time, making it suitable for various IoT applications.

Example: How Netflix utilises Flink for real-time recommendations based on user behaviour

Netflix, the popular streaming service, utilises Flink to process real-time data from user interactions like searches, clicks, and watch history. This enables Netflix to offer personalised recommendations to users, enhancing their streaming experience.

Elasticsearch: Real-time search and analytics

Elasticsearch, an open source search and analytics engine, is a powerful tool for real-time data analysis. It's specifically built for horizontal scalability, reliability, and real-time search capabilities. In the realm of IoT, where data needs to be searched and analysed in real time, Elasticsearch proves to be an invaluable asset.

Elasticsearch operates on a schema-free JSON document model, making it flexible and adaptable to various data formats. It can handle a massive volume of data and provide near-instantaneous search results, which is crucial for real-time analytics in IoT.

Example: How NASA uses Elasticsearch to monitor space data in real time

NASA, the renowned space agency, employs Elasticsearch to monitor and analyse vast amounts of data coming from satellites and space missions. The real-time capabilities of Elasticsearch enable NASA to swiftly identify anomalies, track satellite positions, and analyse mission data, contributing to more efficient space exploration.

InfluxDB: Time-series data handling

InfluxDB, an open source time-series database, is designed to handle high write and query loads. It's the go-to solution for applications that generate data with timestamps, which is common in IoT. Time-series databases are crucial for real-time analytics, as they allow efficient storage, retrieval, and analysis of time-stamped data.

Example: How IoT-based weather stations use InfluxDB to store and analyse climate data in real time

IoT-based weather stations use InfluxDB to store real-time climate data such as temperature, humidity, and atmospheric pressure. The ability of InfluxDB to handle time-series data efficiently enables meteorologists and researchers to analyse weather trends and make timely predictions.

Grafana: Data visualisation

Grafana, an open source analytics and monitoring platform, is a vital component for real-time data visualisation. It integrates with various data sources, including Elasticsearch and InfluxDB, to create interactive and customisable dashboards, enabling users to visualise data in real time.

Example: How a smart home monitoring system uses Grafana to display real-time energy consumption

In a smart home, IoT devices continuously generate data related to energy consumption. Grafana integrates with IoT platforms to display this data in real time, allowing homeowners to monitor and optimise their energy usage.

Advantages of open source tools in real-time analytics

Real-time data analytics in IoT can be a game-changer, and employing open

source tools in this realm offers a multitude of advantages.

- ***Cost-effectiveness:*** Open source tools are usually free to use, reducing the overall cost of implementing real-time analytics solutions. This is especially advantageous for startups and smaller organisations with budget constraints.
- ***Flexibility and customisation:*** Open source tools provide flexibility to tailor solutions according to specific IoT requirements. Developers can customise the tools to match the unique characteristics of the IoT devices and the data they generate.
- ***Strong community support:*** The open source community is vast and active, offering continuous support, updates, and improvements. This collective knowledge and collaboration drive the rapid evolution and enhancement of open source tools.
- ***Security through transparency:*** Open source tools often have their source code available for inspection. This transparency allows experts to assess the security measures and identify vulnerabilities, resulting in more robust and secure solutions.
- ***Rapid prototyping and development:*** Open source tools come with pre-built modules and frameworks, accelerating the development process. Developers can leverage these resources to rapidly prototype and iterate their real-time analytics solutions.

Challenges and considerations

While open source tools offer significant advantages, implementing real-time data analytics in IoT using these tools is not without challenges. It's essential to be aware of these challenges and carefully consider them during the implementation process.

- ***Scalability challenges:*** Real-time data analytics in IoT involves processing a massive volume of data. As the IoT ecosystem expands, ensuring scalability becomes crucial. Open source tools must be able to scale seamlessly to handle increasing data loads.
- ***Data security and privacy concerns:*** Real-time analytics often involves processing sensitive data. Ensuring the security and privacy of this data is paramount. Implementing adequate security measures and complying with privacy regulations are essential considerations.
- ***Latency:*** Achieving true real-time analytics requires minimising latency in data processing. The architecture and design of the analytics solution, along with the efficiency of the open source tools used, play a vital role in minimising latency.
- ***Interoperability:*** IoT environments consist of diverse devices and platforms. Ensuring seamless interoperability between various IoT devices and open source analytics tools is a challenge that requires careful planning and integration.
- ***Complexity of implementation:*** Implementing a real-time analytics solution using open source tools can be complex, especially for those new to the technology. It's essential to have a skilled development team with expertise in both IoT and the selected open source tools.

Addressing these challenges involves a combination of thorough planning, proper architecture design, and leveraging the features of open source tools. Organisations should conduct a comprehensive analysis of their IoT needs and carefully choose the right open source tools that align with their requirements. Moreover, engaging with the open source community and seeking expert guidance can significantly mitigate these challenges.

Predictions for future trends

The landscape of real-time data analytics in IoT is continuously evolving, and several trends are shaping its future.

- ***Edge computing integration:*** With the rise of edge computing, real-time analytics will increasingly move closer to the data source. Edge analytics will become more powerful, reducing latency, and enhancing the efficiency of real-time data processing.
- ***Machine learning integration:*** Integrating machine learning algorithms into real-time analytics will enable IoT devices to make more intelligent and informed decisions on-the-fly. Predictive analytics and anomaly detection will become integral parts of real-time processing.
- ***Enhanced security measures:*** As IoT devices become more prevalent, ensuring the security of real-time analytics will be paramount. Future trends will focus on enhancing security measures to protect sensitive data processed in real time.
- ***Advancements in 5G technology:*** The rollout of 5G will revolutionise real-time analytics in IoT by significantly improving network speed and reducing latency. This will facilitate faster data processing and more efficient real-time analytics.

Staying informed about the latest trends and advancements in both IoT and open source technologies is crucial. With the right combination of open source tools and a forward-thinking approach, organisations can unlock the true power of real-time data analytics in IoT. END

By: Meghraj Singh Beniwal

The author is the founder of Schooling for Next Generation, a tech enthusiast and Android hobbyist. He is a test automation analyst working in Virginia in the US.

The Future of IoT

IoT (Internet of Things) and IoT devices are evolving to enhance productivity and efficiency in a range of industries, as well as ensure more security of data.



The human lifestyle evolves with each social, economic and technological revolution. In recent times, Industry 4.0 has harnessed technology (including cloud platforms, IoT-connected devices, blockchain-enabled crypto transactions, and digital twins) to enhance mass production. One drawback of Industry 4.0 adoption is its focus on either efficiency or productivity. In an efficiency-driven solution, the aim is cost optimisation, resource management, and waste reduction. In a productivity-driven approach, the focus in on reducing time to market and operational overheads.

Industry 5.0 has emerged to address these shortcomings of Industry 4.0. For example, Industry 4.0 introduced robotic automation for improved efficiency and quality of service. Now, in Industry 5.0, we will employ cobots or collaborative robots designed to work closely with humans.

Industry 5.0 also aims not to replace human efforts but to support them, focusing on the quality of service along with cost, resource, and effort optimisation. It seeks to improve customer relations with contact center AI, interactive bots, mixed reality-based immersive customer experiences, and a responsive and distributed supply chain for seamless collaboration with teams.

However, the transition to Industry 5.0 is challenging and demands advancements in IoT, connected communication, cloud-based scalable business solutions, IoT-enabled data transmission and management, hyper automation for operational maturity models, and immersive user interactions.

Business context of IoT services

When the Internet of Things (IoT) is applied to different sectors, new terms and definitions emerge, such as IoMT, IoDT, IoVT, IoIT and IoRT. Each of

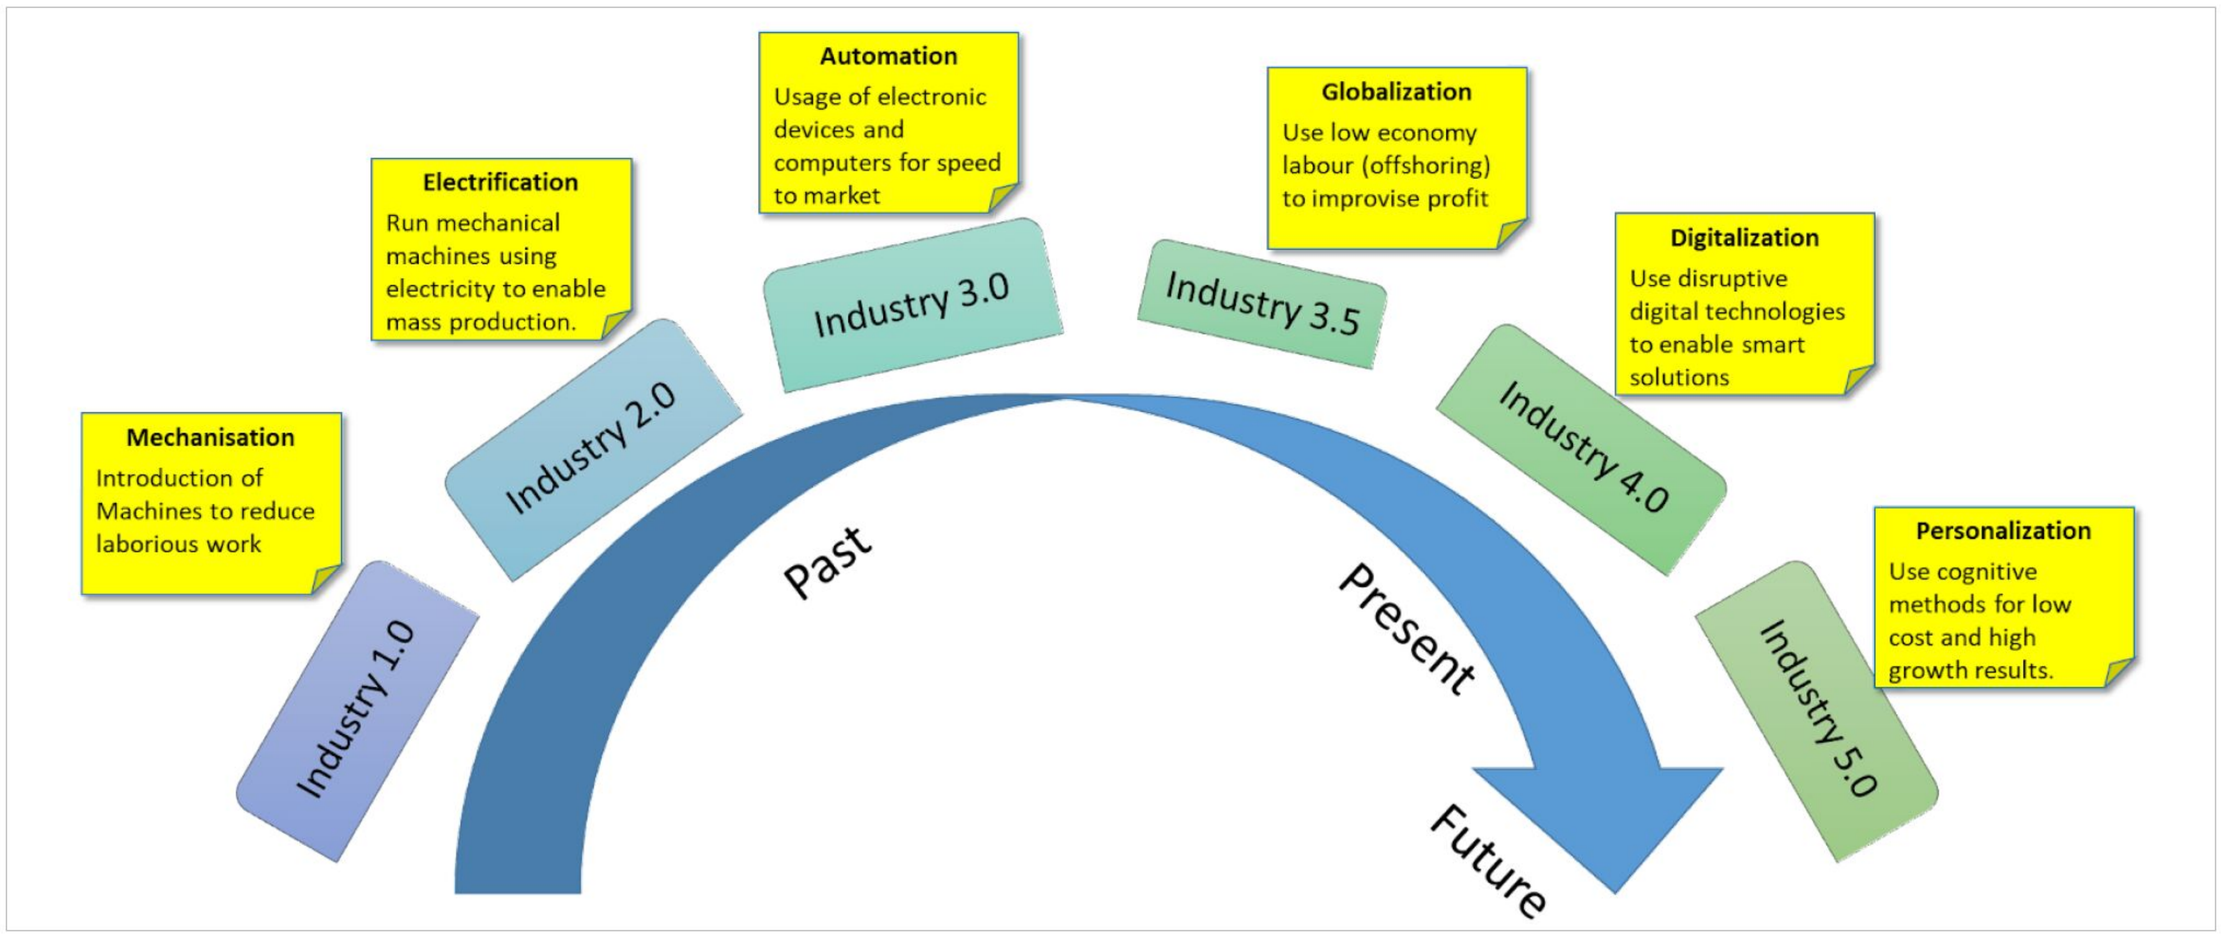


Figure 1: Evolution cycle of Industry 5.0

these simplifies some complex business and technical requirements. Let's explore these terms.

- ***Internet of Medical Things (IoMT) –*** Commonly used for technology solutions combined with sensors for health monitoring and remote patient monitoring in the healthcare industry
- ***Internet of Drone Things (IDT)*** – Utilises drones and unmanned aerial vehicles (UAV) for applications like agriculture, wildlife monitoring, evacuation drills, and industrial site monitoring
- ***Internet of Digital Twins (IoDT)*** – Replicates human-like virtual assistants, robotic assistants and chatbots to aid end users, e.g., virtual banking assistants for financial advisors
- ***Internet of Vehicle Things (IoVT)*** – Applied in the automotive industry for digital metaspace to sell cars, simulators for vehicle driving training, and vehicle design using digital canvases
- ***Internet of Robotic Things (IoRT)*** – Used in production design to create a low-cost simulated environment that uses robots for inspection, quality tests and training
- ***Internet of Intelligent Things (IoIT)*** – Applied in industries like virtual/online retail stores to analyse user patterns and behaviour to provide intelligent services based on machine learning

Internet of Everything (IoE) combines all or some of these services to integrate various processes, using the internet for data handling, as well as specialised devices like head-mounted devices, IoT sensors and monitors, and people with defined roles/responsibilities to handle various business functions.

In simple terms, these services (IoE) are not confined to specific technologies but use various combinations to enhance customer experiences, provide better solutions, and reduce costs for 'go to market' solutions across different sectors.

Trends in IoT-enabled application architecture

The latest trends in IoT-enabled application architecture include:

- ***Blockchain for IoT security:*** Blockchain technology can secure IoT devices and data by providing a decentralised and tamper-proof system.
- ***Smart home devices:*** IoT-enabled smart home devices are gaining popularity, offering remote control and communication with other household devices.
- ***Cybersecurity priority:*** With the growing prevalence of IoT devices, cybersecurity becomes increasingly important. IoT-driven software architectures must consider the implications of software complexity, security, and privacy.
- ***Edge computing rising over cloud computing:*** Edge computing is gaining prominence due to its ability to process data faster, reducing the need to transmit data to the cloud for processing.
- ***AI and Big Data significance:*** IoT is a key source of real-time data for AI applications, enabling AI-driven decisions.
- ***Complex wearable devices:*** IoT-enabled wearable devices used for health monitoring and fitness tracking are becoming more complex.
- ***Microservices architecture adoption:*** IoT-centric software architecture poses unique challenges, and adopting microservices architecture can help.

Overall, IoT-enabled application architecture must evolve to meet the unique demands and challenges posed by IoT, as emerging IoT devices become the new norm.

Security requirements in IoT-enabled application services

IoT-enabled devices communicate through protocols like Message Queue Telemetry Transport (MQTT), which is an open messaging and lightweight protocol for IoT. This involves personal area networks (PANs) or wide area networks (WANs) connecting multiple data centres through public or private cloud services.

In this network, IoT security on infrastructure components and IoT application security on IoT-enabled edge services (like secure access service edge (SASE)) are important to protect connected devices and application data. Some of the core principles for developing a security strategy for IoT applications and infrastructure services are:

- Data encryption at rest to protect data when it is stored in the data storage or application storage facility. This could be through encryption or device level cipher protection to avoid tampering data access from outside.
- Data encryption in transit or in motion to protect data when transporting through a communication protocol. This could be through secured tunnel or safeguarding network services using secured protocols like SSL, TLS, FTPS, etc.
- Device identity to protect unauthorised devices being connected to IoT network of services. Each device in the IoT implementation is identified through unique device identity. This is used to authenticate (verify device access) and authorise (permit device connection) in the IoT ecosystem of connected devices.
- Protect device authentication through a token-based authentication and authorisation framework like OAuth 2.0.
- Policy management to provide role-based access control, along with user principles and policy definitions to enable device protection facilities in the IoT ecosystem. This will enable access management for all the connected services.
- Two-way protection in all services through certificate based authentication like public key infrastructure (PKI) and organisation-based access control (OrBAC).

When using MQTT for machine-to-machine network communication, it's essential to ensure device authorisation and data protection. IoT security and IoT application security play critical roles in the IoT ecosystem.

Edge computing for IoT services

Edge computing for IoT services is a distributed computing paradigm that brings computation and data storage closer to data sources, such as IoT devices and sensors. This allows for real-time data processing and analysis, reduced latency, and improved bandwidth efficiency.

Edge computing is particularly well-suited for IoT applications because it can help to address some of the key challenges of IoT, such as:

- ***Handling large volumes of data:*** IoT devices generate large volumes of data, which can be costly and time-consuming to transmit to the cloud. Edge computing can help reduce this amount of data by processing and filtering it at the edge.
- ***Reducing latency:*** IoT applications often require real-time data processing and analysis in order to make timely decisions. Edge computing can help to reduce latency by bringing computation closer to the sources of data.
- ***Improving reliability:*** IoT devices are often deployed in remote or harsh environments where network connectivity may be unreliable. Edge computing can help to improve reliability by allowing IoT devices to continue operating even when they are not connected to the cloud.

Edge computing for IoT services can be used to implement a wide variety of applications, including:

- ***Real-time monitoring and control:*** Edge computing can be used to monitor and control IoT devices in real time. This can be used for applications such as predictive maintenance, traffic management, and industrial automation.
- ***Video analytics:*** Edge computing can be used to perform video analytics on IoT devices, such as security cameras and traffic cameras. This can be used for applications like object detection, motion detection, and facial recognition.
- ***Machine learning:*** Edge computing can be used to train and deploy machine learning models on IoT devices. This allows for real-time predictions and decisions to be made at the edge, without the need to send data to the cloud.

Here are some examples of edge computing for IoT services in use today.

- ***Smart cities:*** Edge computing is being used to power smart city applications such as real-time traffic monitoring, parking management, and air quality monitoring.
- ***Industrial automation:*** Edge computing is being used to improve the efficiency and reliability of industrial automation processes. For example, edge computing can be used to monitor and control machines in real time, and to perform predictive maintenance to prevent downtime.
- ***Retail:*** Edge computing is being used to improve the customer experience in retail stores. For example, it can be used to power personalised recommendations, in-store navigation, and queue management systems.

Overall, edge computing is a powerful technology that can be used to improve the performance, reliability, and security of IoT services.

Threats in IoT enabled architecture

As per the Singapore Computer Society (SCS) Cyber Literacy Report 2020 and Blackberry Threat Report 2022, IoT attacks remain one of the top 5 cyber threats around the world. IoT plays a vital role in today's connected world, spanning applications from smartwatches and IP cameras in traffic and road safety control, to CCTVs for smart city administration, PoS monitoring in retail stores, and inventory supply chain in manufacturing sector.

IoT architecture is based on three major components.

- ***Edge client –*** Collects data from connected devices and sends it to the gateway
- ***Gateway*** – Validates and forwards data to the cloud platform for storage and processing
- ***Cloud*** – Receives and processes data blocks from the gateway, stores them in cloud storage, and transfers for processing like analytics and reporting services hosted on the cloud native platform

Security design is vital in IoT architecture design strategy. We need to see this at four levels:

- ***Infrastructure security*** to manage security at the physical device level from device controllers, edge, gateway and cloud.
- ***Platform security*** to manage security at runtime services in edge and gateway in order to handle access control and protect data privacy requirements.
- ***Network security*** to manage, monitor and adopt data encryption during network transfers — two-way handshake to ensure secured clients are connected, and audit security logs for anomaly behaviour, if any.
- ***Application security*** to manage native and third party services handling data and to ensure that IoT services like API are protected from cyberattacks.

Potential challenges in IoT security could be data privacy compromised when using third party services, unprotected data storage, lack of role-based access control (RBAC), weak cloud architecture to protect public exposed services, and insecure or outdated IoT device components being exposed to cyberattacks.

Future developments in IoT enabled services

Here are some potential future developments in IoT-enabled services.

- ***Ubiquitous connectivity:*** 5G and beyond will enable massive scale IoT deployment. Connectivity will be faster, cheaper, and more reliable.
- ***AI at the edge:*** More advanced edge devices will allow localised AI inferencing, reducing latency and data transfer.
- ***Increased adoption in industries:*** Manufacturing, healthcare, smart cities and more will see expanded IoT adoption for monitoring, automation, etc.
- ***Advanced analytics:*** With more data, analytics and predictive modelling will become more sophisticated and accurate.
- ***Growth of as-a-Service models:*** Device management, connectivity, security, etc, will increasingly be offered 'as a service'.
- ***New voice/gesture interfaces:*** There will be more natural interfaces based on voice assistants, computer vision, and AR/VR.
- ***Increased focus on security:*** Security architectures spanning edge, network, and cloud will be crucial.
- ***Distributed computing:*** Swarms of collaborative edge devices for distributed computing paradigms will be leveraged.
- ***Personalised IoT:*** Tailored solutions using wearables, smart home tech, and personalised healthcare will be popular.
- ***Sustainability:*** Greener devices, renewable energy, and optimised energy usage will be in focus.
- ***Smart cities:*** Urban IoT rollouts will improve traffic, public services, and more.
- ***Better hardware:*** Smaller, cheaper, more efficient IoT sensors and edge hardware are on the cards.
- ***Closed loop automation:*** Systems that translate IoT data into automated intelligent actions to improve business processes will be in common use.

The future of IoT promises a brighter, more efficient, and interconnected world, but it will require overcoming scale, complexity, privacy and security challenges. END



References

- *https://nexusintegra.io/industry-5-0-the-new-revolution/*
- *https://www.techtarget.com/iotagenda/tip/Top-8-IoT-applications-and-examples-in-business*
- *https://www.interact-lighting.com/global/iot-insights/the-iot-and-new-business-models*
- *https://www.iotworldtoday.com/connectivity/iot-trends-2021-iot-accelerates-digitization-automation-and-edge-architecture*
- *https://www.checkpoint.com/fr/cyber-hub/network-security/what-is-iot/iot-security-architecture/*
- *https://www.embitel.com/blog/embedded-blog/iot-cloud-and-application-security*

By: Dr Magesh Kasthuri and Dr Anand Nayyar

Dr Magesh Kasthuri is a senior distinguished member of the technical staff and principal consultant at Wipro Ltd. This article expresses his views and not that of Wipro.

Dr Anand Nayyar is a PhD in wireless sensor networks and swarm intelligence. He works at Duy Tan University, Vietnam, and loves to explore open source technologies, IoT, cloud computing, deep learning and cyber security.

Generative AI:
The Things It Can Help With

Generative AI amplifies tech investments, including AI models. It is here to stay, and will aid businesses perform better as it continues to evolve. Almost every industry can benefit from it.



Today, every IT vendor is promoting generative AI. In current seminars and conferences on IT technology, leaders invariably talk of 'generative AI'. McKinsey defines generative AI as algorithms (such as ChatGPT) that can be used to create new content, including audio, code, images, text, simulations and videos. Generative AI is primed to make a significant impact on enterprises over the next five years.

Key characteristics of generative AI technology

Generative AI comprises models that describe our requirements, visualising and generating content to match prompts. It expedites ideation, brings vision to life and allows more time for creativity. It can create a wide variety of data such as images, videos, audio, text and 3D models, often employing large language models, like ChatGPT.

The following are the key characteristics of generative AI technology.

- ***Text management:*** Generative AI can complete a given text coherently, translate languages, summarise text concisely, and generate text that mimics human writing.
- ***Contextual understanding:*** Generative AI has a strong ability to understand the context.

- By 2026, generative design AI will automate 60% of the design effort for new websites and mobile apps. – *Gartner*
- By 2027, nearly 15% of new applications will be automatically generated by AI without human involvement. This is not happening today. – *Gartner*
- Recent CEO surveys show almost 80% of CEOs believe AI is likely to significantly enhance business efficiency in their organisations. – *Forbes*
- Today, an estimated 60% of IT leaders are looking to implement generative AI. – *CIO.com*
- 44% of the worker's core skills are expected to change in the next five years. Training employees to leverage generative AI is going to be critical. – *World Economic Forum*

- ***Natural language processing (NLP):*** It can perform various NLP tasks, processing human language and allowing users to ask questions conversationally.
- ***Answering questions:*** Generative AI can answer questions based on a knowledge base.
- ***Personalisation:*** It can be fine-tuned for specific use cases.
- ***Multi-language support:*** Generative AI can handle NLP tasks in multiple languages.
- ***Advanced semantic search:*** It enables semantic search in large volumes of structured and unstructured data, including databases, documents, etc. It also supports auto data indexing using custom AI/ML based data processing for efficient embedding.
- ***Content moderation:*** It has smart content filtering and moderation engines for query and response, trainable on enterprise-specific data.
- ***Integration:*** It can integrate with various applications and data sources within an organisation, including CRM, ERP, and other proprietary systems, to access and analyse data via APIs.
- ***Role-based access control:*** The system can be configured with role-based access control, ensuring that users only have authorised access to data.
- ***Software coding:*** Generative AI can generate, translate, explain and verify code.

Real-world use cases of generative AI

Generative AI's use cases are vast and continually evolving. Businesses across industries are experimenting with different ways to incorporate generative AI to enhance efficiency and decision-making capabilities. Generative AI applications improve experiences, reduce costs and increase revenues for enterprises.

Here is a summary of the generative AI use cases across various industries.

Healthcare and pharma

Generative AI-based applications help healthcare professionals be more productive by identifying potential issues upfront, providing insights to deliver interconnected healthcare and improve the health of patients. It helps in:

Better customer experience: Automates administrative tasks such as processing claims, scheduling appointments, and managing medical records.

Patient health summary: Helps provide healthcare decision support by generating personalised patient health summaries, speeding up patient response times, and improving the patient experience.

Faster analysis of publications: Generative AI helps in reducing the time it takes to create research publications on specific drugs by analysing vast amounts of data from multiple sources faster than ever.

Personalised medicine: Generative AI-based individualised treatment plans are tailored to a patient's genetic makeup, medical history, life style, etc.

Healthcare virtual assistant: It provides end users with conversational and engaging access to the most relevant and accurate healthcare services and information.

Manufacturing

Generative AI enables manufacturers to make better use of their data, leading to advancements in predictive maintenance and demand forecasting. It also helps in simulating manufacturing quality, improving production speed, and material efficiency.

Predictive maintenance: It helps in estimating the life of machines and their components. Generative AI can provide proactive information to technicians about repairs and replacement of parts and machines, thus reducing downtime.

Performance efficiency: Anticipating the problems proactively can cover production disruptions, bottlenecks, and safety risks in real time.

Other uses in the manufacturing industry include:

- Yield, energy and throughput optimisation
- Digital simulations
- Sales and demand forecasting
- Logistics network optimisation

Retail

Generative AI in retail helps in personalising offerings, brand management, as well as optimising marketing and sales activities.

Personalised offerings: It enables retailers to deliver customised experiences, offerings and pricing, and plan according to customer demand, modernising the online and physical buying experience.

Dynamic pricing and planning: It helps predict demand for different products, enabling more confident pricing and stocking decisions.

Other uses include:

- Campaign management
- Content management
- Augmented customer support
- Search engine optimisation

Banking

Generative AI applications help in delivering a personalised banking experience to customers. There are other benefits too.

Risk mitigation and portfolio optimisation: Generative AI helps banks to build data foundations for developing risk models, and identify events that are impacting the bank. It can help mitigate risks to optimise portfolios.

Customer pattern analysis: Generative AI can analyse patterns in banking data at scale, helping relationship managers and customer representatives to identify customer preferences, anticipate needs, and create personalised banking experiences.

Customer financial planning: Generative AI can be used to automate customer service, identify trends in customer behaviour, and predict customer

needs and preferences. This helps to understand the customer better and provide personalised advice.

Generative AI also helps with:

- Anti-money laundering regulations
- Compliance
- Financial simulations
- Applicant simulations
- Risk analytics
- Fraud prevention

Insurance

Generative AI's capability of analysing and processing large amounts of data helps in accurate risk assessments and effective claims processes.

Customer support: Generative AI can provide multilingual customer service by translating customer queries and responding to them in the preferred language.

Policy management: It analyses large amounts of unstructured data related to customer policies, various policy documents, customer feedback, and social media literature to implement better policy management.

Claims management: Generative AI helps in analysing various claims artefacts to enhance the overall efficiency and effectiveness of claims management.

It also helps with:

- Customer communications
- Coverage explanations
- Cross-selling and upselling of products
- Acceleration of the product development life cycle
- Innovation of products

Education

Generative AI helps to connect teachers and students. It enables collaboration between teachers, administrators and technology innovators to provide better education.

Student enablement: Generative AI helps students with real-time lesson translations in different languages. It helps visually impaired students with classroom accessibility.

Student success: Deep analytic insights into student success help teachers make informed decisions on how to improve performances.

Telecommunications

Generative AI adoption by the telecom industry improves operational efficiency and network performance. Here, generative AI can be used to:

- Analyse customer purchasing patterns
- Give personalised service recommendations
- Enhance sales
- Manage customer loyalty
- Get insights into customer preferences
- Provide better data and network security, enhancing fraud detection

Public sector

The goal of digital government initiatives is to establish a connected government and provide better citizen services.

Smart cities: Generative AI helps in toll management, traffic optimisation, and sustainability.

Better citizen services: It provides citizens with easier access to connected government services through tracking, search, and conversational bots.

It also enables:

- Service operations optimisation
- Contact centre automation

Benefits of generative AI

The many benefits of generative AI include:

- Increased productivity
- Better content creation
- Personalised customer experiences
- Identification of new customer journeys and audiences
- Improved customer interactions through enhanced chat and search experiences
- Enhanced creativity and data exploration
- Exploration of large amounts of unstructured data through conversational interfaces and summarisations
- Transformation of campaigns, audiences, experiences, and insights

There is a significant gap between what generative AI can do and what enterprises want it to do. From assessing business needs to building generative AI solutions, it is necessary to make the right choice for enterprise business use cases.

However, generative AI is here to stay and will continue to improve in future versions. Unlike previous technologies, AI can make increasingly complex decisions enabling new business opportunities. The use of generative AI across enterprises is becoming more and more widespread, possibly even trending towards industrialisation.

Generative AI is an opportunity and not the competition. It won't replace humans but will assist in the business success of the next generation. END

Acknowledgements

The author would like to thank Santosh Shinde of BTIS, Enterprise Architecture division of HCL Technologies Ltd, for the support in bringing this article to fruition as part of architecture practice efforts.

By: Dr Gopala Krishna Behara and Meghna Behara

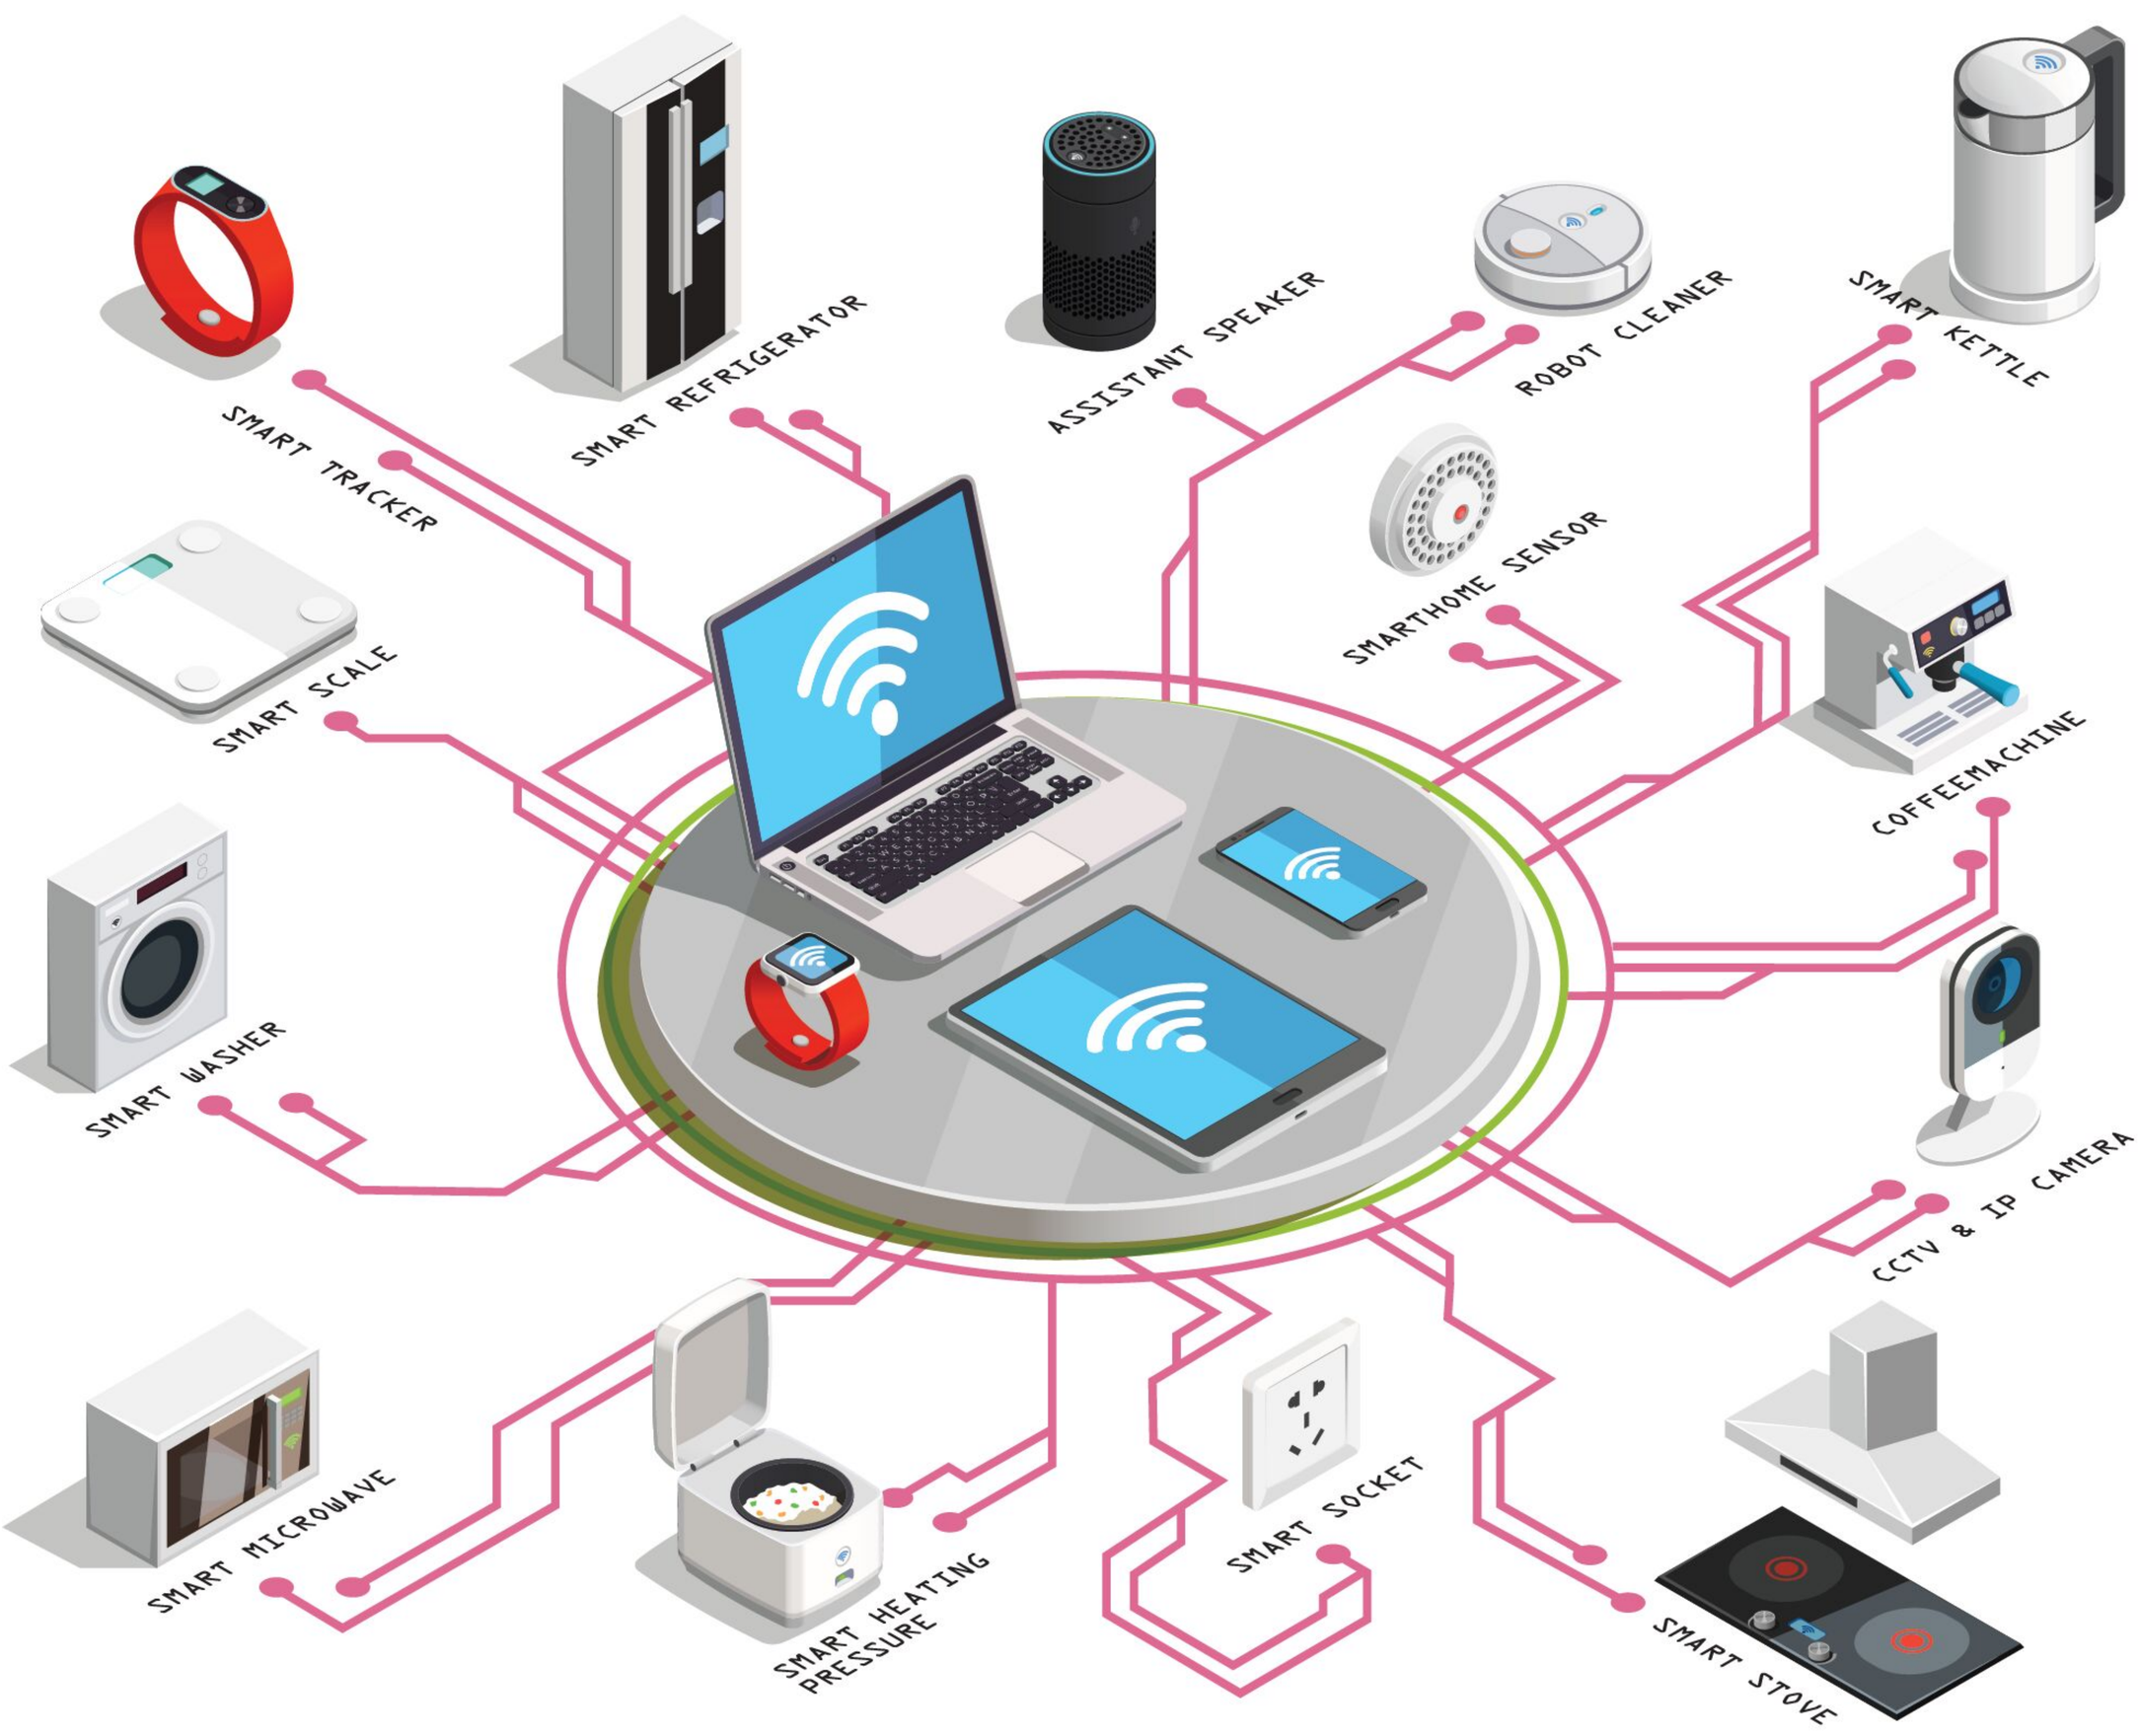
Dr Gopala Krishna Behara is an enterprise architect in the BTIS Enterprise Architecture division of HCL Technologies Ltd. He has a total of 28 years of experience in the IT industry.

Meghna Behara is an artificial intelligence enthusiast, and is pursuing her specialisation in the subject.

Disclaimer: The views expressed in this article are that of the authors and HCL does not subscribe to the substance, veracity, or truthfulness of the said opinion.

IoTivity: Using Open Source to Make Devices Talk to Each Other

With a burgeoning ecosystem of stakeholders, IoTivity is poised to shape the future of IoT, fostering innovation and enhancing the lives of individuals and organisations around the globe. What makes it special is a deep commitment to open source principles.



The Internet of Things (IoT) has been reshaping our world, bringing convenience and connectivity to every corner of our lives. We've witnessed smart homes that adjust the temperature before we even think of it and industrial setups that run like clockwork, thanks to IoT. But amid all the marvels, a significant challenge looms: how do we get these devices to talk to each other seamlessly, regardless of who made them or what language they speak? That's where the Open Connectivity Foundation (OCF), now IoTivity, comes in.

IoTivity is an open source software project enabling seamless device-to-device connectivity. It helps billions of wired and wireless IoT devices securely connect to each other and to the internet. The Open Connectivity Foundation (OCF) develops specifications, interoperability guidelines, and a certification program for these devices.

Let's discover the story of IoTivity, its mission, and why it matters in the IoT universe.

The birth of IoTivity

IoTivity's journey began back in 2014 when the Open Interconnect Consortium (OIC) was born. Its mission was clear: establish open standards for IoT device communication and interoperability. Two years later, in 2016, the OIC underwent a transformation and rebranded itself as the Open Connectivity Foundation (OCF), broadening its horizons to encompass more than just connectivity.

At its core, IoTivity was driven by a singular purpose -- creating a world where IoT devices could communicate effortlessly, transcending boundaries set by manufacturers or communication protocols. The vision was to ensure that smart devices collaborated harmoniously, offering enhanced user experiences while maintaining stringent security.

Open source, open possibilities

The key to IoTivity's success lies in its unwavering commitment to open source principles. IoTivity provides an open source software framework that acts as a blueprint for implementing their standards. This open source approach is a game-changer. It fosters collaboration, empowering developers and organisations to embrace these standards without being shackled by proprietary solutions.

IoTivity's open source ethos encourages innovation. It paves the way for startups and tech giants alike to join hands in refining and expanding IoT standards. With a common framework and protocols, IoTivity quickens the pace of IoT solution development, ensuring compatibility across different ecosystems.

Interoperability: The heartbeat of IoTivity

At the heart of IoTivity's mission lies interoperability — the ability of devices from diverse origins to communicate seamlessly. IoT encompasses an array of devices, each using distinct communication protocols and technologies. Without interoperability, these devices remain isolated silos, limiting their utility.

IoTivity tackles this challenge head-on by defining open standards that promote interoperability. These standards cover a wide spectrum of IoT aspects, including device discovery, data modelling, communication protocols, and security. By adhering to these standards, IoT device manufacturers ensure that their products can effortlessly communicate with devices from other manufacturers, creating a united IoT ecosystem.

Prioritising IoT security

Security is paramount in the IoT realm. The proliferation of connected devices creates new avenues for attacks, making it imperative to safeguard sensitive data and maintain system integrity. IoTivity takes security seriously, developing features and protocols to ensure secure device discovery and communication.

IoTivity's security framework encompasses robust mechanisms for device authentication, secure data exchange, and access control. By incorporating these security features, IoTivity mitigates the risks associated with IoT, instilling trust among users and organisations embracing IoT solutions.

Standardisation for harmony

One of IoTivity's most impactful contributions to the IoT landscape is the establishment of standardised protocols and specifications. These standards provide a common language for IoT devices and services, facilitating effective communication.

IoTivity leverages well-established internet protocols such as CoAP (Constrained Application Protocol) and HTTP for device communication. Moreover, it offers a framework for managing device resources and data models. These standardised elements simplify the lives of developers,

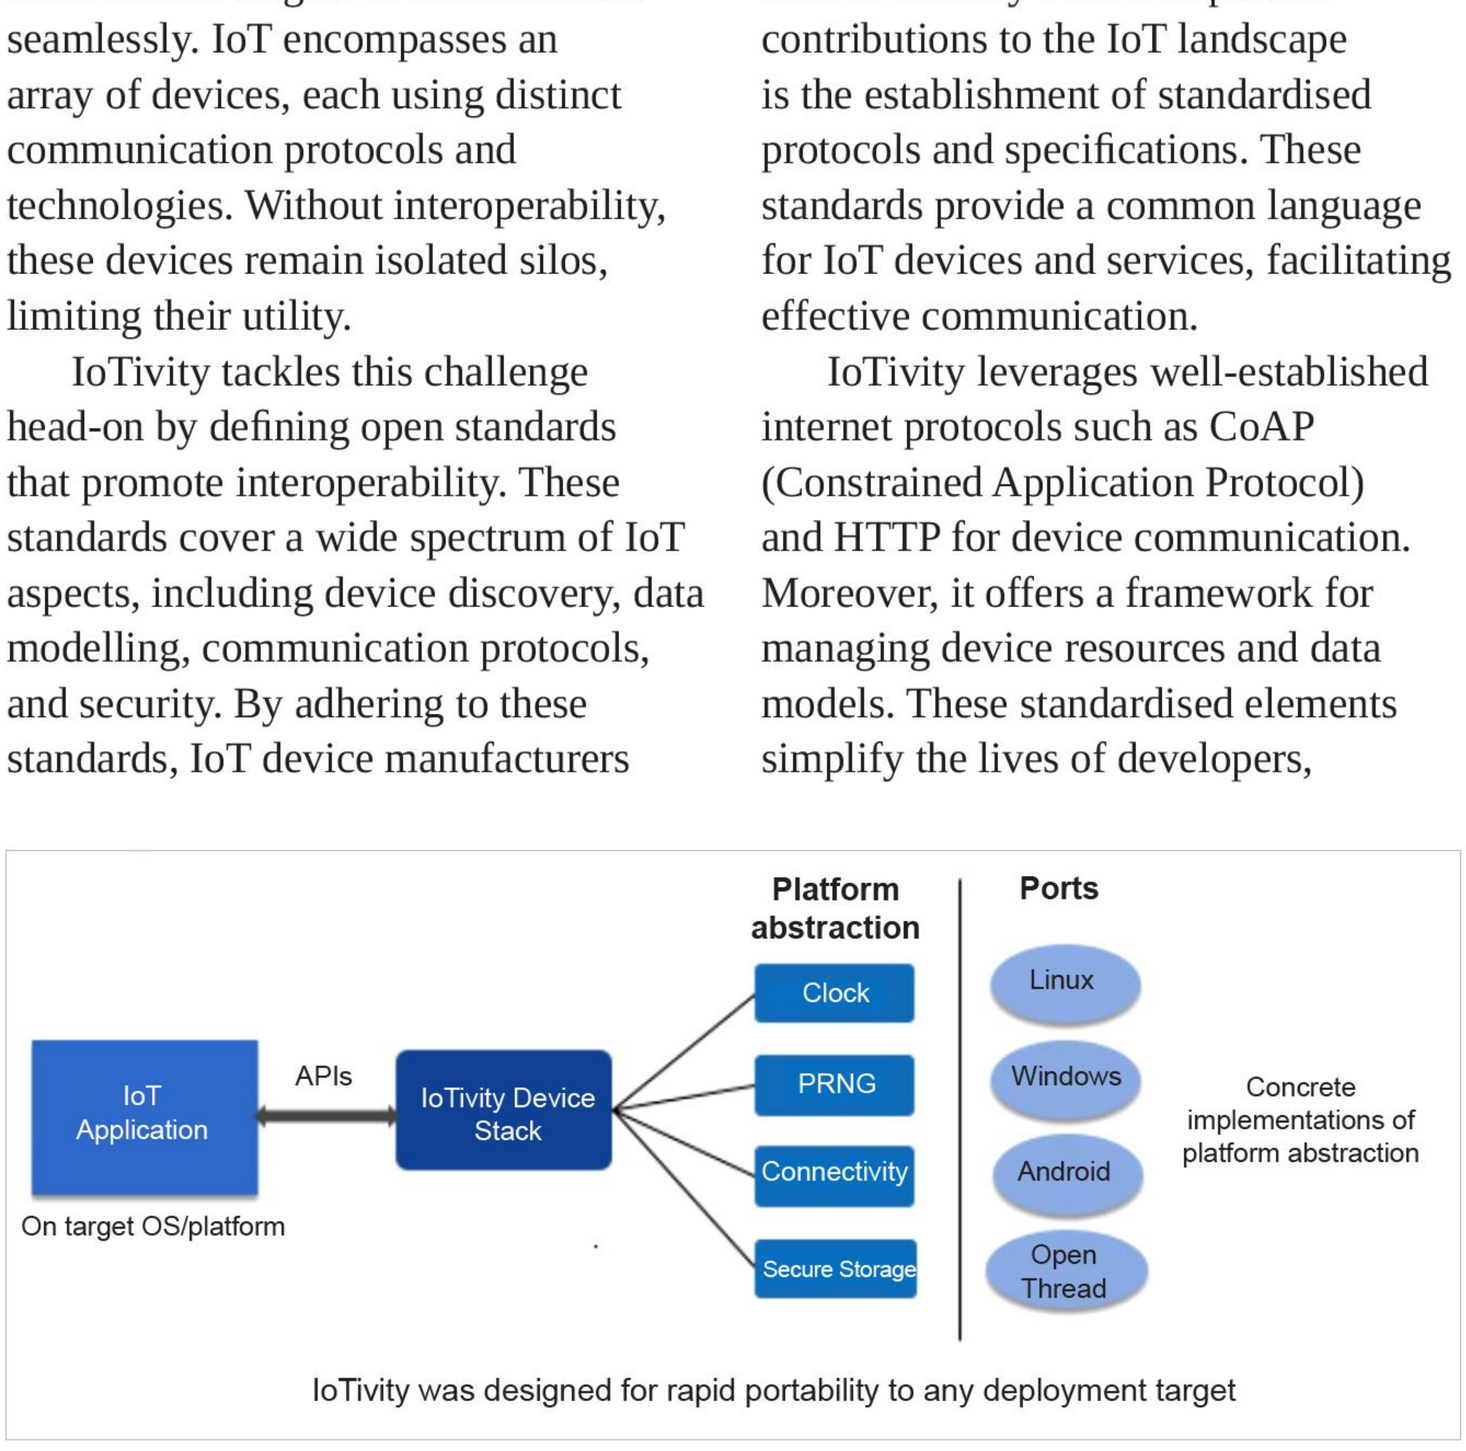


Figure 1: IoTivity architecture

enabling them to create IoT applications and services that work harmoniously with a wide range of devices.

A thriving ecosystem of innovation

IoTivity has cultivated a thriving ecosystem of member organisations. This ecosystem encompasses industry leaders, device manufacturers, and stakeholders deeply entrenched in the IoT landscape. These members unite to co-create and champion IoTivity's standards and frameworks.

The diversity within this ecosystem ensures that IoTivity's standards draw from a broad spectrum of perspectives and use cases. Furthermore, it promotes the adoption of these standards across various industries, spanning smart homes, healthcare, automotive, and industrial automation.

Real-world impact

IoTivity's influence extends beyond theoretical standards and frameworks. Its practical applications manifest in real-world IoT projects and products. A shining example is the realm of smart homes.

IoTivity's standards have revolutionised the smart home domain, enabling interoperability between devices from different manufacturers. This means our smart thermostat can harmoniously communicate with our smart lights, regardless of their respective brands. Such interoperability enhances user convenience and accelerates the adoption of smart home technologies.

In the healthcare sector, IoTivity has played a pivotal role in the development of connected medical devices that share data securely and reliably. This has the potential to elevate patient care, facilitate remote monitoring, and streamline healthcare operations.

Current state of IoTivity

IoTivity's influence is palpable in diverse domains. We delve into the applications and sectors where it has made the most impact.

Smart homes: IoTivity's standards have revolutionised the smart home domain, enabling interoperability between devices from different manufacturers. This means smart thermostats can harmoniously communicate with smart lights, regardless of their respective brands. Such interoperability enhances user convenience and accelerates the adoption of smart home technologies.

Healthcare*:* In the healthcare sector, IoTivity has played a pivotal role in the development of connected medical devices that share data securely and reliably. This has the potential to elevate patient care, facilitate remote monitoring, and streamline healthcare operations.

Industrial automation*:* IoTivity's standards have found applications in industrial automation, where seamless communication between diverse machines and sensors is critical for optimising operations and ensuring safety.

Why to use IoTivity

IoTivity offers a host of compelling reasons to enhance your IoT endeavours. The most prominent ones are briefed below.

- ***Interoperability:*** By adopting IoTivity, we can ensure that our IoT devices can communicate seamlessly with other devices, regardless of their origins, making it easier to create comprehensive and integrated IoT solutions.
- ***Open source:*** IoTivity is an open source project, making its codebase freely available to developers and organisations. This open source nature encourages a broader community of contributors, lowers development costs, and speeds up the development of IoT solutions.
- ***Security:*** IoTivity places a strong emphasis on security. It includes features and protocols for secure device discovery, communication, and data exchange. By using IoTivity, we can leverage these built-in security features to protect our IoT ecosystem and sensitive data.
- ***Resource management:*** IoTivity provides a framework for managing device resources and data models. This simplifies the creation of IoT applications and services by providing a standardised way to interact with and manage IoT devices.
- ***Future-proofing****:* By adopting IoTivity, we future-proof our IoT solutions. As IoT evolves and new devices and technologies emerge, having a standardised and interoperable foundation ensures that our IoT ecosystem can adapt and grow without major disruptions.

The journey of the Open Connectivity Foundation, now known as IoTivity, stands as a testament to the power of open standards and collaboration in advancing the IoT landscape. Through its open source framework, the development of open standards, and a steadfast commitment to security, IoTivity has laid the foundation for a more interconnected and secure IoT ecosystem.

As the IoT world continues to evolve and expand, IoTivity's contributions will remain instrumental in realising the vision of a seamlessly connected world. END

References

- *http://iotivity.org/*
- *https://openconnectivity.org/technology/iotivity/*

By: Neetesh Mehrotra

The author is working in EY as a senior. His areas of interest are Java development and automation testing.

Neuromorphic Computing: The Future of IoT

Neuromorphic computing transcends mere connectivity, creating a future where our tech mirrors human cognition. Get ready to embrace brain-inspired tech!



IoT is a rapidly evolving field. The recent developments in AI, and now, quantum computing, have accelerated this process by a factor of 100. From smart homes to interconnected urban infrastructures, the IoT ecosystem thrives on continuous data exchange and instant processing. Yet, as vast as these networks become, there's an inherent need to reimagine how we perceive and approach computing for them.

Enter neuromorphic computing—a revolutionary concept that takes inspiration from the most intricate and efficient computer known to us: the human brain.

At its core, neuromorphic computing seeks not only to replicate the structural intricacies of our neural networks but also to emulate the adaptability, efficiency, and real-time processing capabilities of the brain.In doing so, it promises to overcome many of the limitations that traditional computing architectures face when applied to the sprawling IoT networks of the future.

In this article, we'll journey through the foundational architecture that has been the backbone of our computer systems—the von Neumann architecture. We delve into the evolving realm of neuromorphic computing, and explore why this shift is the next big step in the evolution of IoT.

The traditional von Neumann architecture in IoT

The von Neumann architecture has long been the cornerstone of modern computing. Named after the renowned mathematician and physicist John von Neumann, this design has dictated how computers process information for decades. Central to this architecture is the principle of storing both data and instructions (programs) in the same memory space and processing them sequentially. Figure 1 explores the von Neumann architecture.

The key components of von Neumann architecture are:

- ***Central processing unit (CPU):*** The brain of the computer, responsible for executing instructions
- ***Memory:*** Stores both the data and the program instructions
- ***Input/output (I/O) systems:*** Devices and interfaces that allow the computer to interact with the external world, like sensors in IoT devices

- ***Bus system:*** Communication pathways that transport data between different components

In the context of IoT, von Neumann's architecture has played a pivotal role. IoT devices, at their heart, are miniaturised computers that interact with the physical world. They collect data through sensors, process this data, and then take actions or communicate findings. This sequence of operations—input (sensing), processing, and output (acting or communicating)—is a direct manifestation of von Neumann's principles.

However, this architecture does come with its set of limitations.

- ***Bottleneck:*** Often termed the 'von Neumann bottleneck', there's an inherent limitation in the speed at which data can be transferred between the CPU and memory. In real-time environments like IoT, this can cause latency issues.
- ***Power consumption:*** von Neumann architectures can be power-hungry, not ideal for IoT devices which often rely on battery power or energy harvesting.
- ***Parallel processing:*** The sequential nature of this architecture isn't always best suited for tasks that demand parallel processing, a feature that's becoming increasingly essential with the influx of large-scale data in IoT.

Despite these challenges, von Neumann's architecture has been the foundation of many breakthroughs in IoT. But as the scale and complexity of IoT systems grow, and the demand for real-time, efficient processing increases, there's a pressing need for architectures that can better handle these requirements. This sets the stage for the emergence of neuromorphic computing.

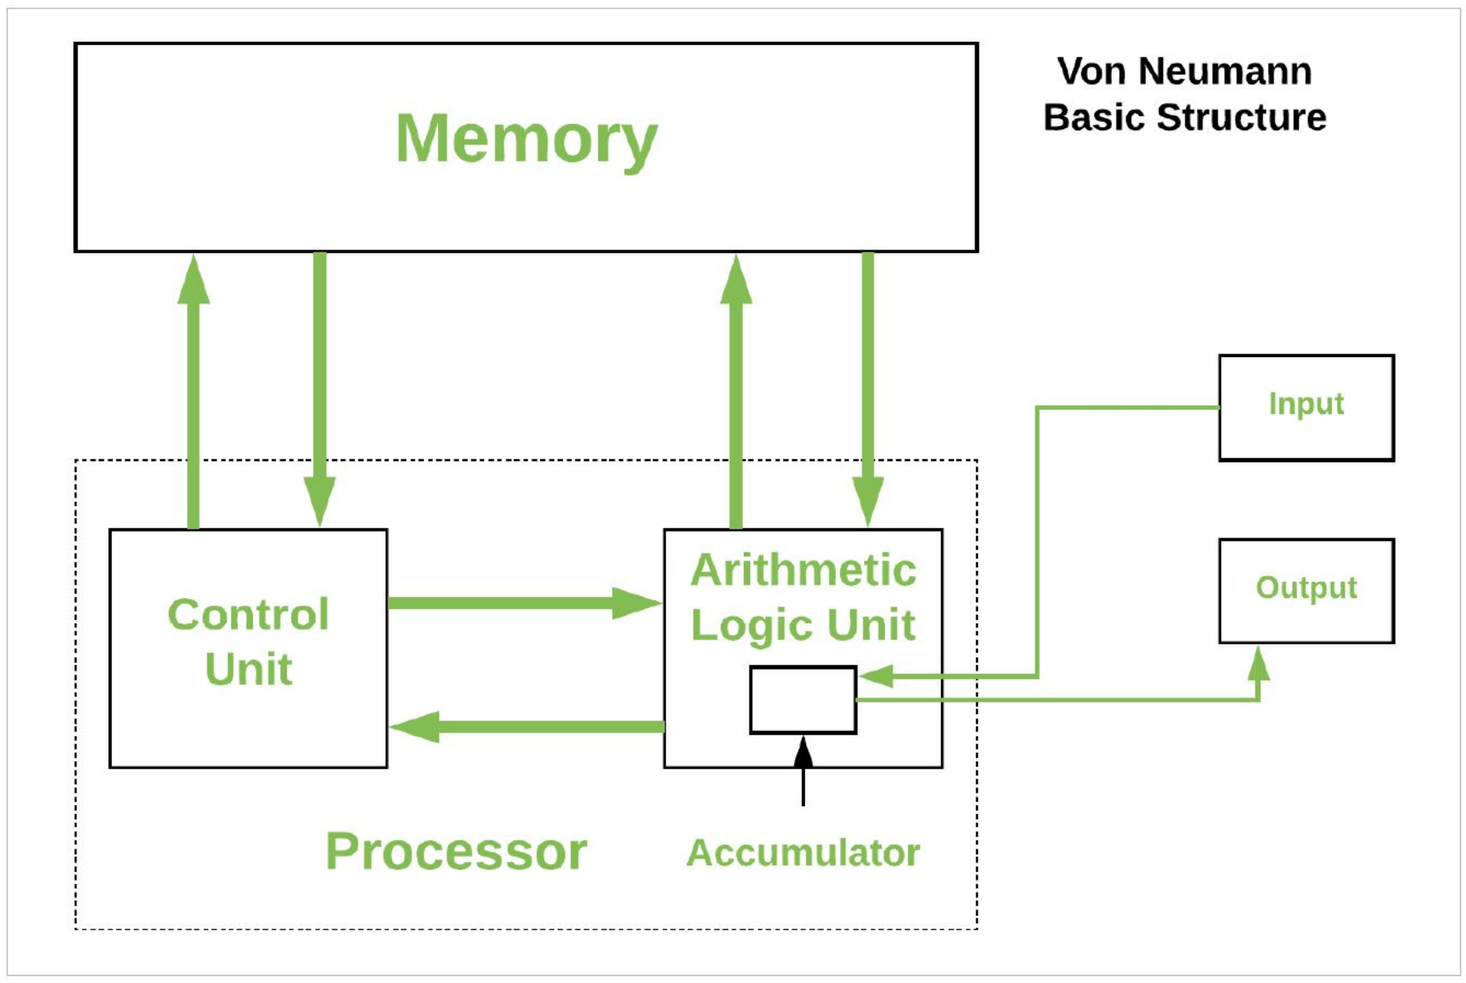


Figure 1: von Neumann architecture (Google Images)

Transition to neuromorphic computing

In order to transcend the constraints of traditional computing systems, researchers have turned to an unlikely yet profound source of inspiration: the human brain. The brain, with its billions of neurons and trillions of synaptic connections, operates in ways that no man-made machine can replicate — at least not yet. This realisation birthed the field of neuromorphic engineering, a domain that seeks to design algorithms, hardware, and systems inspired by the structure and functionality of biological neural systems. Figure 2 explores the architectural overview of neuromorphic computing.

Principles of neuromorphic computing

***Parallelism over sequence*:** Unlike von Neumann architectures, which work predominantly in sequences, neuromorphic systems employ massively parallel operations, emulating the brain's vast network of simultaneously active neurons.

Event-driven processing: Neuromorphic systems often use event-driven models, such as spiking neural networks (SNNs). Instead of continuous computations, these models rely on discrete events or 'spikes', mirroring how neurons either fire or remain inactive.

Adaptability and learning: In traditional architectures, hardware and software are distinct entities. In contrast, neuromorphic systems blur these lines, allowing the hardware itself to adapt and learn, akin to the plasticity of biological synapses.

Memristors: The synaptic mimic

One of the groundbreaking components in neuromorphic engineering is the memristor. Memristors, whose resistance varies based on the history of currents passed through them, are seen as ideal components to represent the synaptic weights in artificial neural networks. Their capability to retain memory even when powered off mirrors the long-term potentiation and depression in biological synapses.

Why the shift matters for IoT

With billions of devices projected to be part of the IoT ecosystem, the demand is clear: low power consumption, real-time processing, and adaptive learning capabilities. Neuromorphic computing, with its energy efficiency and real-time adaptability, offers a promising avenue to address these needs. Devices can potentially learn and

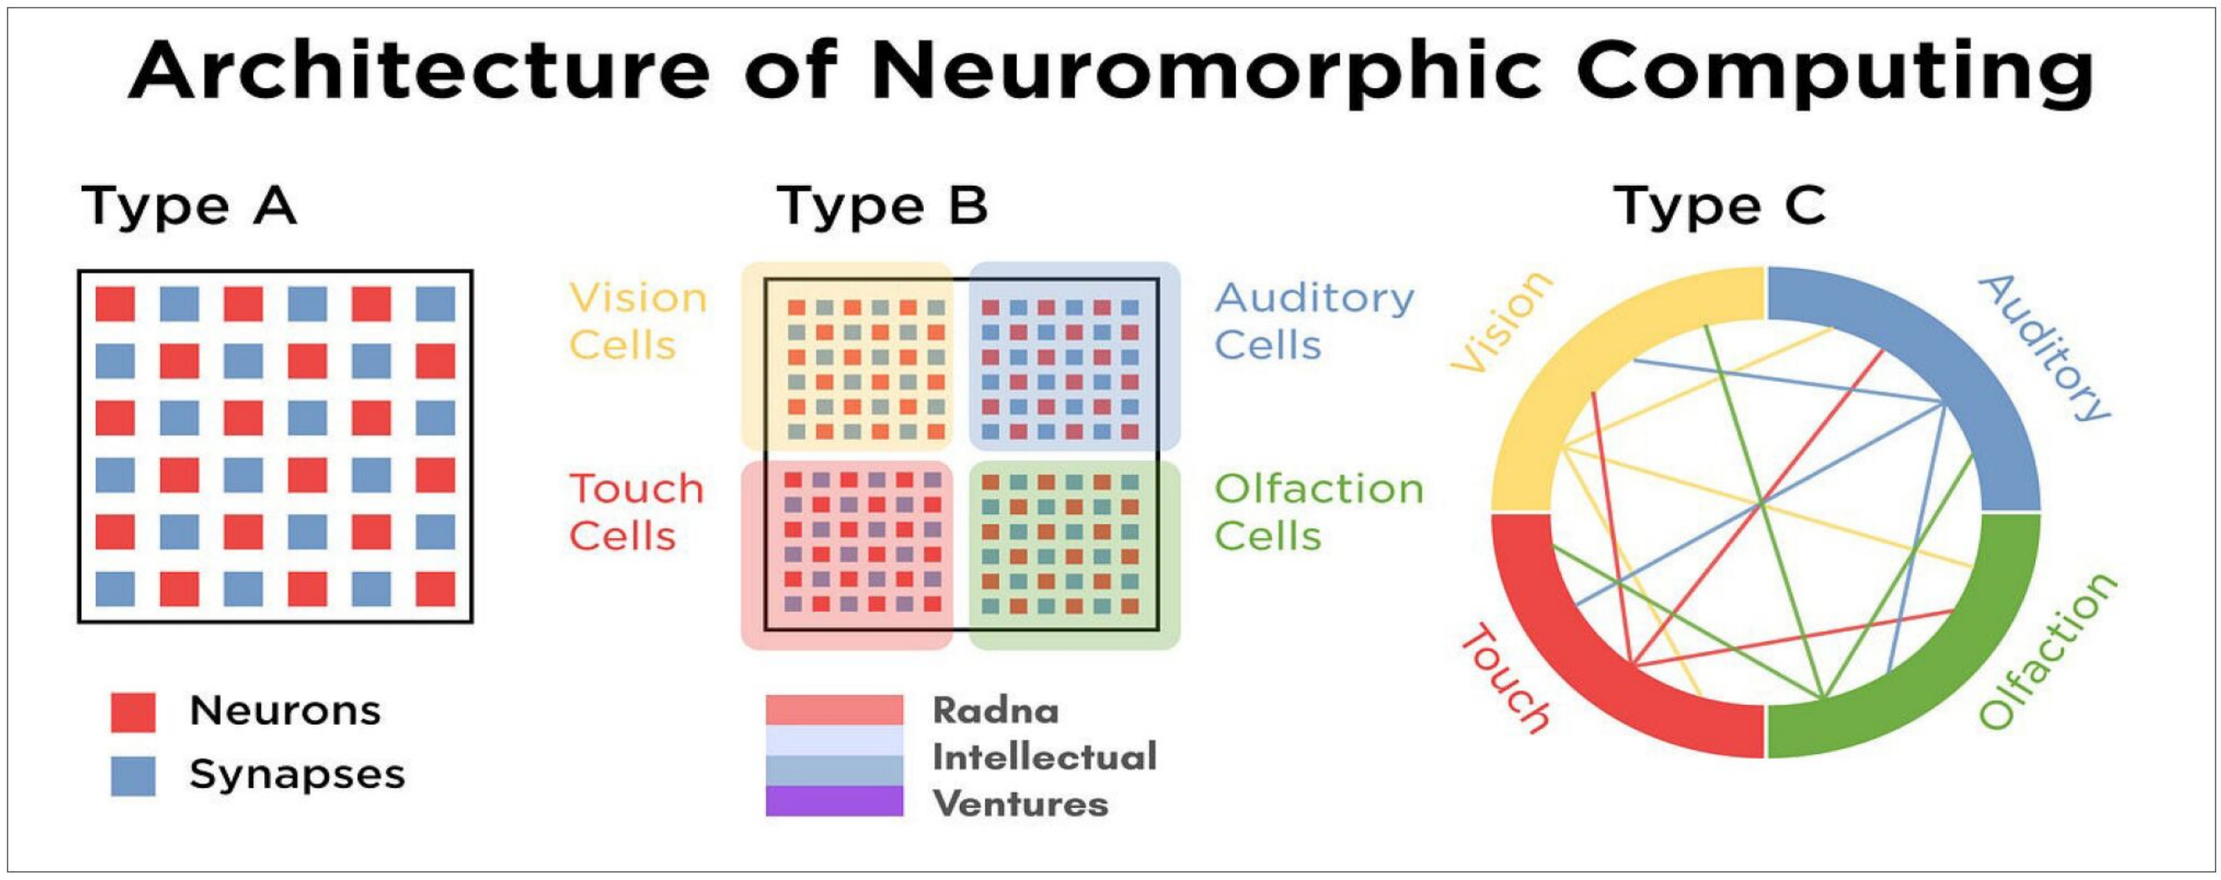


Figure 2: The architectural overview of neuromorphic computing (Google Images)

adapt on-the-fly, make real-time decisions at the edge without relying on centralised servers, and do so with remarkable power efficiency.

In essence, neuromorphic computing isn't just an incremental improvement over existing architectures; it represents a fundamental paradigm shift in how we approach computing for complex, interconnected systems like IoT.

Cross-comparing von Neumann and neuromorphic architectures

The architectural design behind any computing system not only defines its operational methodology but also its potential and limitations. So let us cross-compare the von Neumann architecture with the evolving architecture of neuromorphic computing. Figure 3 provides a cross-comparison between these architectures.

- ***von Neumann architecture***

Sequential processing: At its core, the von Neumann model is built around the fetch-decode-execute cycle. Instructions and data, both stored in memory, are fetched and processed sequentially, leading to inherent delays.

Shared memory bottleneck: The shared memory for both data and instructions often becomes a bottleneck, especially when dealing with complex computations. This is the infamous 'von Neumann bottleneck' that can limit system performance.

Stateful design: von Neumann systems operate in a stateful manner—operations depend on a sequence of previously executed instructions, leading to potential inefficiencies in certain real-time or parallel processing tasks.

- ***Neuromorphic architecture***

Parallelism at heart: Instead of a linear sequence, neuromorphic designs embrace massive parallelism, with artificial neurons operating concurrently. This mirrors the human brain, where billions of neurons can fire or stay silent simultaneously, allowing for multitasking and rapid sensory processing.

Event-driven modality: Neuromorphic systems, especially those using spiking neural networks (SNNs), operate based on events. A neuron either spikes (fires) or doesn't, making the system inherently efficient as operations are based on events rather than continuous cycles.

Inherent adaptability: By using components like memristors that can adjust their resistance based on historical currents, neuromorphic designs allow for an adaptability that's hardwired into the system. This is akin to how biological synapses can strengthen or weaken over time, allowing for learning and memory.

Why it matters for IoT

The transition from von Neumann to neuromorphic isn't just about mimicking the brain but about addressing real-world challenges. As IoT devices proliferate, the demands for low latency, real-time processing, and adaptive learning are increasing. The latency inherent in the von Neumann design, especially in large-scale data processing scenarios, can be a limiting factor for IoT devices needing real-time responses. Neuromorphic designs, with their parallelism and event-driven nature, can address these challenges, offering a more responsive and adaptable architecture for the next generation of IoT solutions.

Industry examples

As academia has unveiled the principles of neuromorphic computing, industry giants and innovative startups have taken the lead in actualising these concepts.

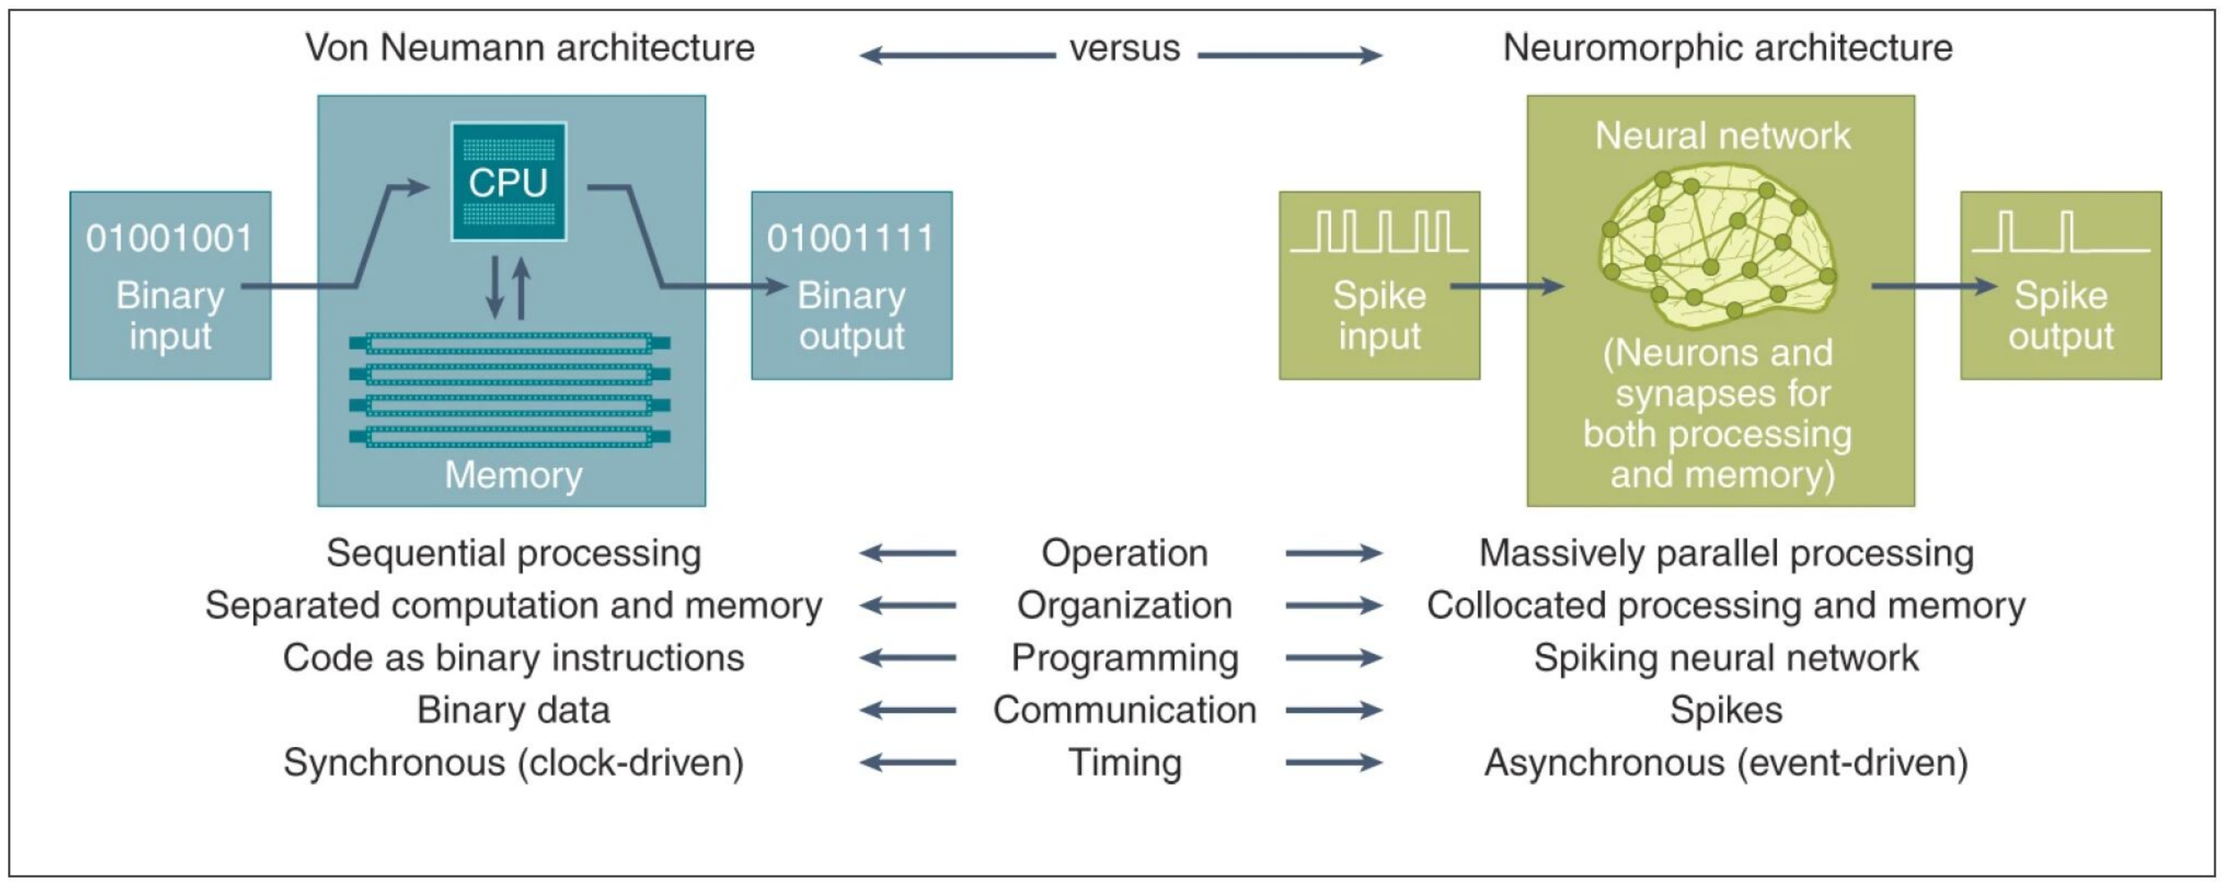


Figure 3: von Neumann vs neuromorphic architectures (*Image credits:* Schuman, C.D., Kulkarni, S.R., Parsa, M. et al. Opportunities for neuromorphic computing algorithms and applications. Nat Comput Sci 2, 10–19 (2022)).

IBM's TrueNorth

Launched by IBM in collaboration with DARPA's SyNAPSE program, TrueNorth stands as one of the pioneering neuromorphic chips.

Specifications: Comprising 5.4 billion transistors, the chip simulates over a million programmable neurons and 256 million programmable synapses, offering a snapshot of the brain's complexity.

Applications: Beyond mere proof-of-concept, TrueNorth has been deployed in real-world scenarios, including real-time hand gesture recognition, optical character recognition, and even powering a prototype mobile eye-tracking system.

Intel's Loihi

Intel's foray into neuromorphic engineering, Loihi is a testament to the industry's commitment to the paradigm. Designed for research purposes, it has been instrumental in fostering exploration and innovation in neuromorphic applications.

Specifications: The chip simulates roughly 8 million neurons, and is designed to provide high efficiency and scalability.

Potential: With on-chip learning, Loihi represents a significant leap towards systems that can adapt and evolve. Intel's research has showcased its capabilities in tasks ranging from path planning for robots to adaptive, real-time control systems for prosthetic limbs.

Neuralink

While not a neuromorphic company *per se*, Elon Musk's Neuralink is weaving the narrative of blending advanced computing with neurobiology. Aimed at creating brain-machine interfaces, Neuralink's endeavours underscore the potential of merging the biological with the artificial.

Relevance: Neuralink's work is emblematic of the broader aspiration within neuromorphic computing: understanding, emulating, and eventually synergising with neural systems. As these interfaces evolve, they might incorporate neuromorphic principles to improve efficiency, adaptability, and biocompatibility.

The vision: Beyond medical applications, Neuralink's long-term vision of symbiosis between the human brain and AI systems aligns with the essence of neuromorphic computing – building machines that resonate with the nuances of human cognition.

Why neuromorphic computing is the future of IoT

The future needs devices that not only communicate but also think, adapt, and evolve. Currently, there are two technologies that are powering this vision into reality, i.e., artificial intelligence (AI) and quantum computing.

Role of AI

- ***Demand for real-time processing:*** AI-driven applications in IoT, from voice assistants to anomaly detection in industrial setups, necessitate real-time data processing. Traditional architectures often grapple with latencies, whereas neuromorphic systems, with their parallelism and event-driven nature, are primed for instantaneous responses.
- ***Adaptive learning on the edge:*** IoT devices often operate at the edge – away from central servers or clouds. With neuromorphic chips, these devices can learn and adapt on-site without constantly needing data transfers to central servers, ensuring efficiency and data privacy.

Continued to page ...68

Smart Healthcare: How IoT is Reshaping the Industry

Healthcare is being transformed by the Internet of Things (IoT), and becoming more effective, data-driven, and patient-centred. We take a look at the transformative impact of IoT on the healthcare industry, outlining its potential to improve patient treatment and increase overall healthcare.



The healthcare sector has seen a significant surge in the adoption of Internet of Things (IoT) devices during 2023**.** Based on a report published by *MarketsandMarkets*, the worldwide market for IoT in the healthcare sector is anticipated to reach $188 billion by 2025**.** This growth is expected to be achieved through a compound yearly growth rate (CAGR) of 29.9% from 2020 to 2025.

Figure 1 shows IoT's growth trajectory in the healthcare business.

Key applications of IoT in healthcare

IoT has diverse applications in healthcare, driving improvements in patient care, clinical processes, and operational efficiency (Figure 2)**.**

Patient monitoring: Through the Internet of Things, doctors can remotely monitor their patients' vital signs, the progress of long-term illnesses, and post-surgery recovery. They can identify potential issues early, and deliver timely interventions, reducing the likelihood of hospital readmissions.

Clinical decision support: IoT applications provide real-time patient data, empowering healthcare professionals to make well-informed decisions, thus improving patient care and reducing the occurrence of medical errors.

Device management: IoT is used to manage the status of medical equipment within healthcare facilities. This minimises equipment loss, optimises asset utilisation, and ensures the availability of critical resources.

Predictive maintenance: IoT aids in predicting when medical devices require maintenance or replacement, thereby reducing downtime and ensuring the continuous availability of equipment for patient care.

Research and innovation: IoT-enabled devices and wearables are employed in clinical trials to collect participant data, significantly improving the efficiency and accuracy of medical research, and aiding in drug development.

IoT benefits in healthcare

There are various benefits of integrating IoT in healthcare (Figure 3).

Enhanced patient engagement: With the support of the Internet of Things (IoT), medical portals

and mobile apps can actively involve patients in their treatment. Consequently, patients are more likely to participate in their care and report higher satisfaction with their healthcare professionals.

Cost savings and operational efficiency: IoT and telehealth have demonstrated their cost-effectiveness for both patients and healthcare systems. A report by Deloitte estimates that IoT in healthcare has the potential to annually save the industry up to $300 billion through remote patient monitoring, preventive care, and improved medication adherence.

Telemedicine and telehealth services: Telemedicine and IoT infrastructure extend healthcare accessibility to remote or underdeveloped areas. Telehealth services also reduce the necessity for in-person visits, improving overall access to healthcare.

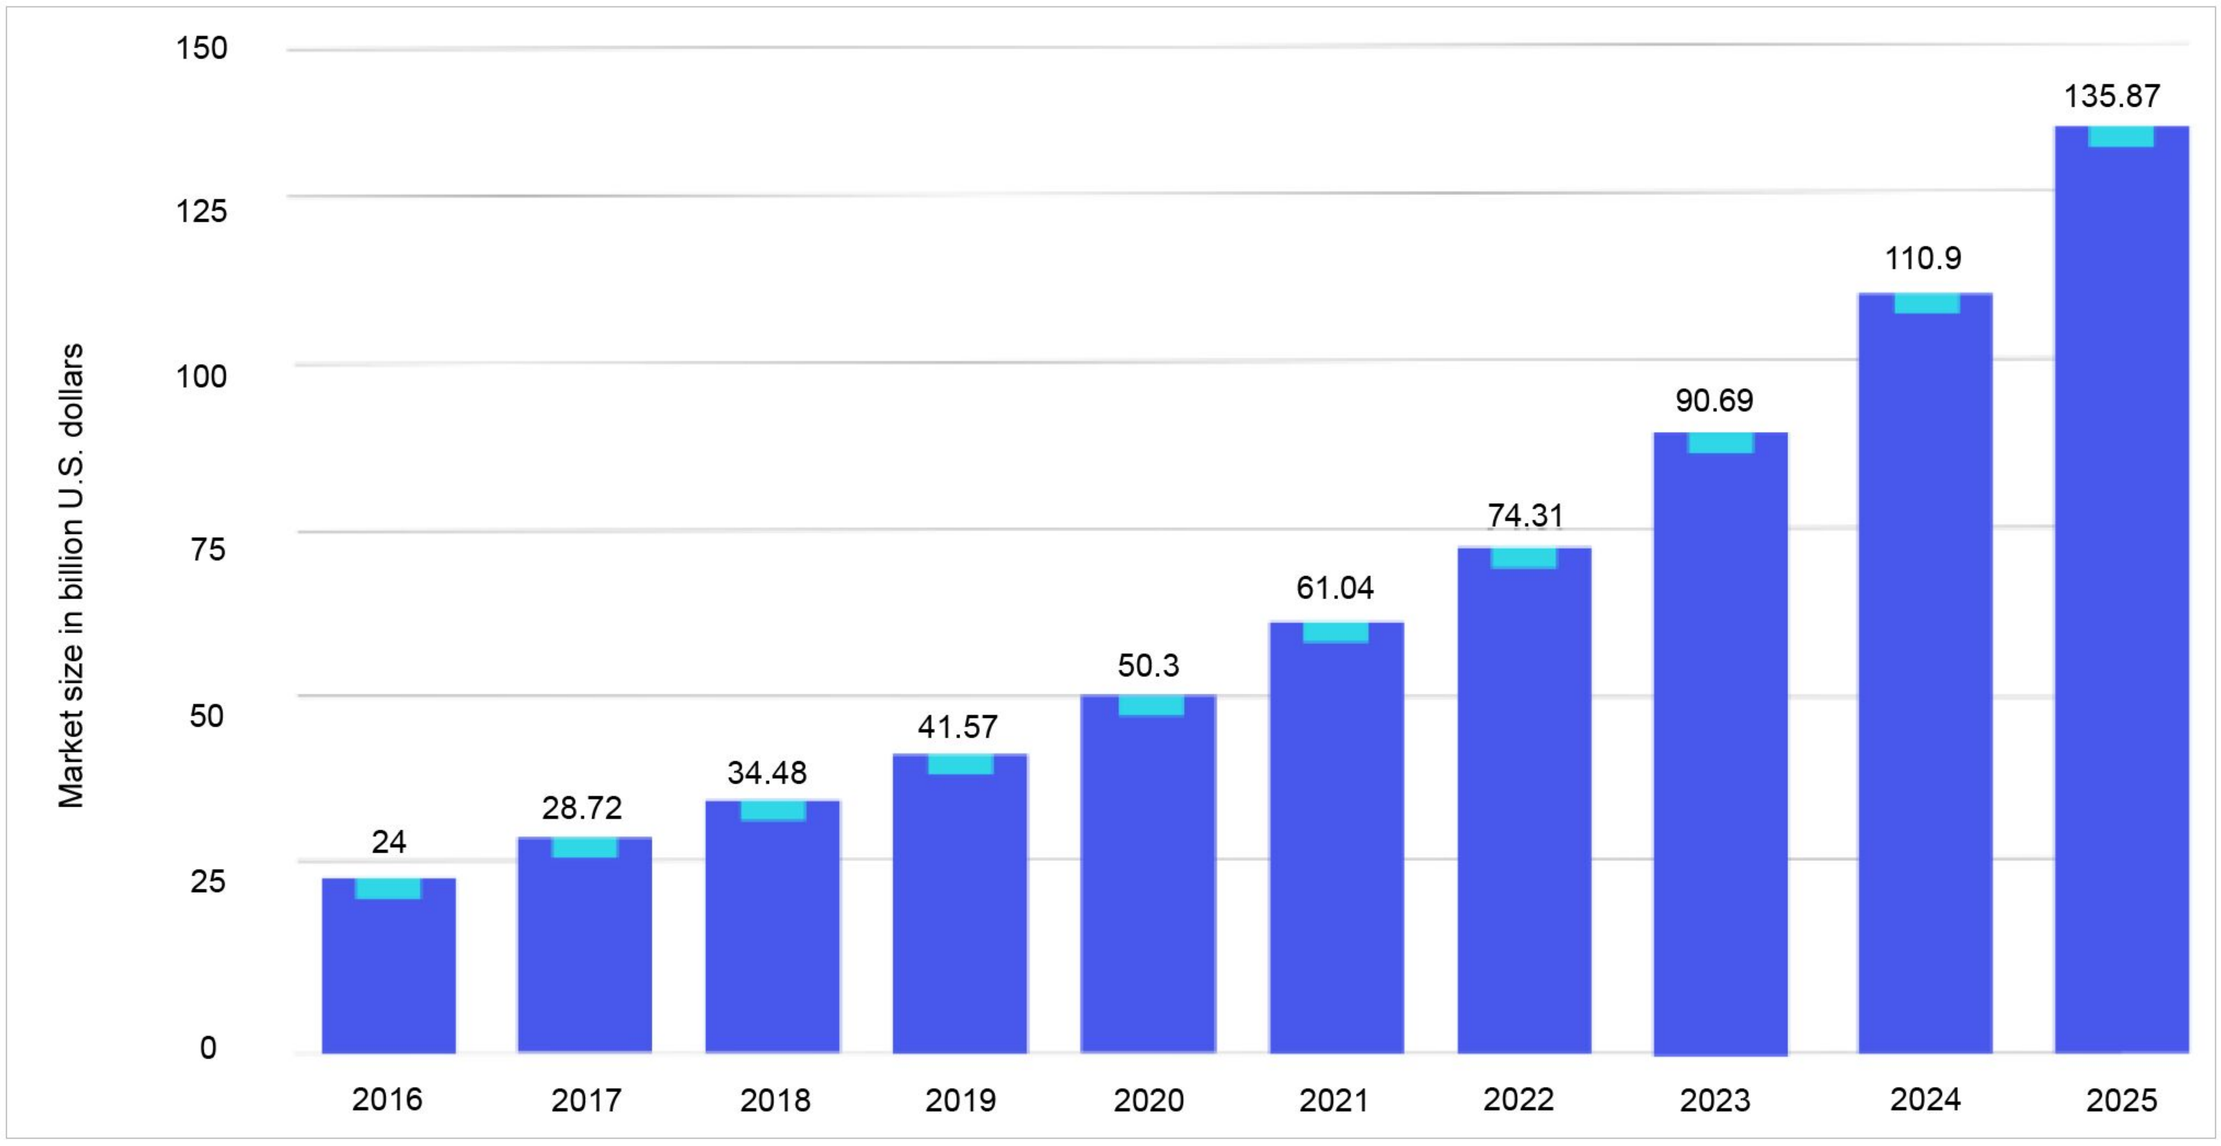


Figure 1: Forecast of healthcare IoT market (2016–2025)

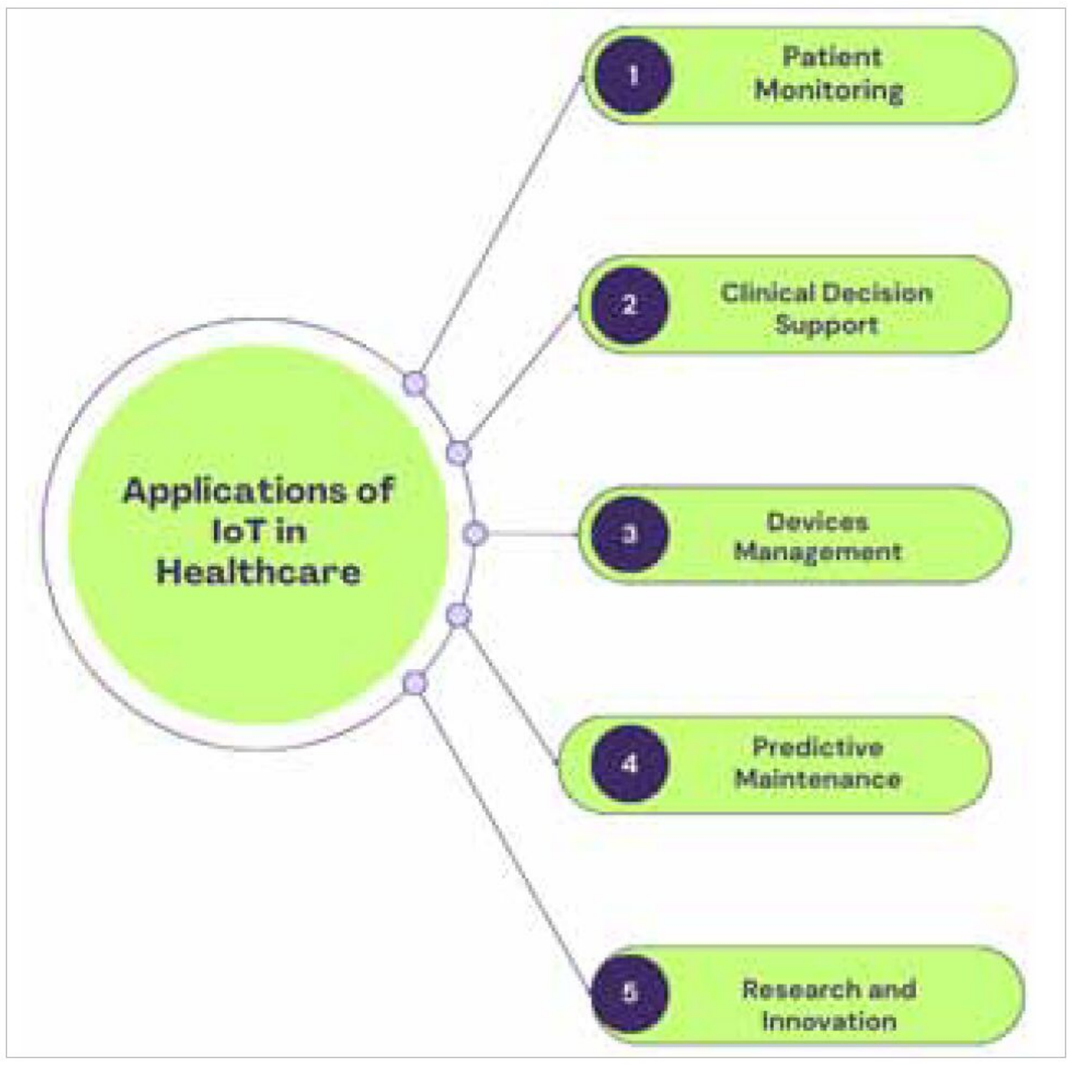


Figure 2: Applications of IoT in healthcare

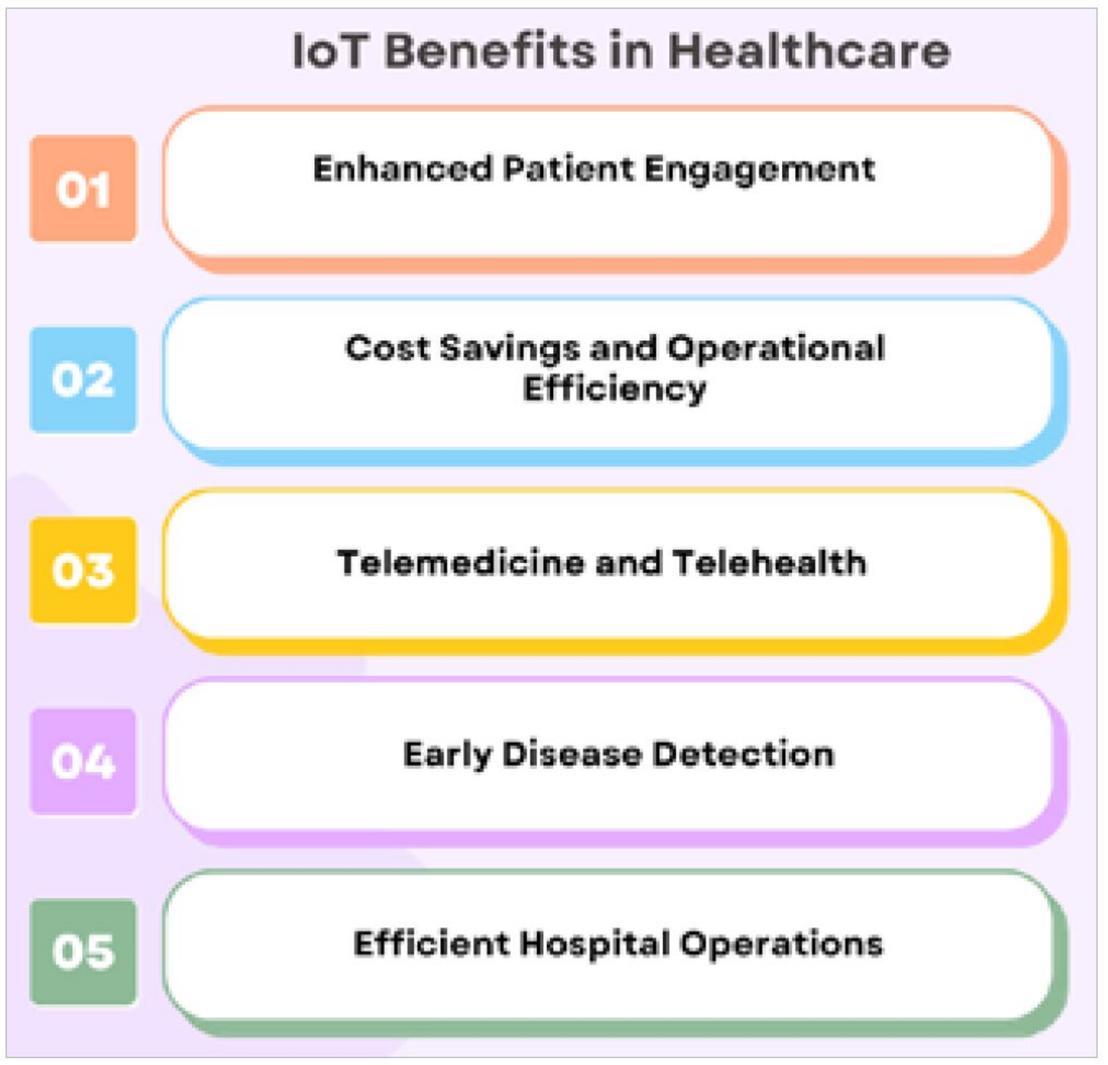


Figure 3: Benefits of integrating IoT in healthcare

Table 1: A comparison of IoT simulation tools

Features	COOJA	NS-3	OMNeT++	NetSim
Simulation purpose	Wireless sensor networks	Network and IoT	Network and IoT	Network and IoT
Licence	BSD	Open source	Academic public licence	Open source
Simulation language	Java	C++ and Python	C, C++	Java, HTML and XML
Mobility support	Yes	Yes	Yes	Yes
Customisation	Limited	Extensible	Highly extensible	Extensible
Network protocols	IEEE 802.15.4, 6LoWPAN	Various IoT protocols	Various IoT protocols	Various IoT protocols
Operating systems	Cross-platform	Cross-platform	Cross-platform	Windows-based
Ease of use	Beginner-friendly	Requires coding skills	Requires coding skills	Beginner-friendly with scripting
Real-time simulation	Limited	Supported	Supported	Limited support
Application areas	WSN and IoT research	A broad range of IoT apps	IoT and communication	IoT research and development

Table 2: Popular Python libraries for IoT-enabled healthcare projects

Library	Description	Features	Use cases	Version
OpenCV	Computer vision and image processing	Image, video processing, object detection, image recognition	Medical imaging, diagnostic analysis, patient monitoring	OpenCV 4.8.0
PyHealth	Machine learning library for healthcare	Simplifies data preprocessing and model deployment	Predictive analytics, patient risk stratification, readmission prediction	PyHealth 1.1.4
MNE-Python	Data analysis for electrophysiology and neuroimaging	Analysis of EEG, MEG, and other neural data, machine learning tools for data interpretation	Neuroimaging research, epilepsy monitoring, brain-computer interfaces	mne 1.5. 1
BlueZ	Bluetooth stack for IoT and healthcare devices	Supports IoT devices' Bluetooth connectivity	Wearable health monitors, Bluetooth-enabled medical devices	BlueZ 5.63
PyDICOM	DICOM file manipulation for medical imaging	Read and write DICOM files – helps extract and manipulate metadata from medical images	Medical image management, radiology, picture archiving, and communication systems (PACS)	PyDICOM 1.1.0
scikit-learn	Machine learning for healthcare data analysis	Supports classification, regression, clustering	Predictive modelling, disease prediction, patient risk assessment	scikit-learn 1.3.1

Early disease detection: IoT devices continuously monitor and analyse patient data in real time, enabling the early detection of signs of diseases or health deterioration. This leads to timely interventions and improved health outcomes.

Efficient hospital operations: IoT optimises hospital operations by tracking equipment, managing inventory, and monitoring environmental conditions. This optimisation leads to cost savings, improved efficiency, and enhanced patient safety.

Open source simulation tools for IoT implementation

Several open source simulation tools can be used for implementing and testing IoT applications. These tools offer flexibility and customisation, making them suitable for many IoT projects. Table 1 gives a comparison of popular IoT simulation tools.

Python libraries for IoT-enabled healthcare implementation

Python is a versatile programming language with many libraries that can be used in IoT-enabled healthcare implementations. Table 2 lists popular Python libraries that are commonly used in IoT-enabled healthcare projects.

IoT is not just changing how we think about healthcare — it's changing how we experience it. Patients are more engaged, doctors better informed, and hospitals vastly more efficient. The Internet of Things is the connective thread, creating a more inter-connected, efficient, and compassionate healthcare system. END

References

- *https://www.marketsandmarkets.com/*
- *https://www.insiderintelligence.com/insights/iot-healthcare/*

By: Dr Aditya Bhardwaj

The author is a B. Tech, M. Tech, and PhD in CSE. He works as an assistant professor in the School of Computer Science Engineering and Technology at Bennett University, Greater Noida. He is experienced in cloud computing and open source technology.

Continued from page ...64

- ***Synergy with deep learning:*** As deep learning models become more intricate, their computational demands skyrocket. Neuromorphic systems, with their inherent efficiency and brain-like processing, could become the preferred substrate for running these complex models, especially in power-constrained IoT environments.

Emergence of quantum computing

- ***Beyond classical computing:*** Quantum computing represents a leap beyond classical computing principles, offering exponentially faster computations for specific tasks. However, the intersection of quantum and neuromorphic computing remains largely uncharted, promising unprecedented computational capabilities.
- ***Potential synergy:*** While neuromorphic systems emulate the brain's structure and functionality, quantum systems could potentially capture its underlying quantum mechanical properties. This synergy could unlock new models of neural processing, further enhancing IoT capabilities.
- ***IoT and quantum:*** As quantum technologies mature, they'll demand architectures that can seamlessly integrate with them. Neuromorphic systems, with their adaptability and bio-inspired designs, may offer the most compatible bridge to this quantum future.

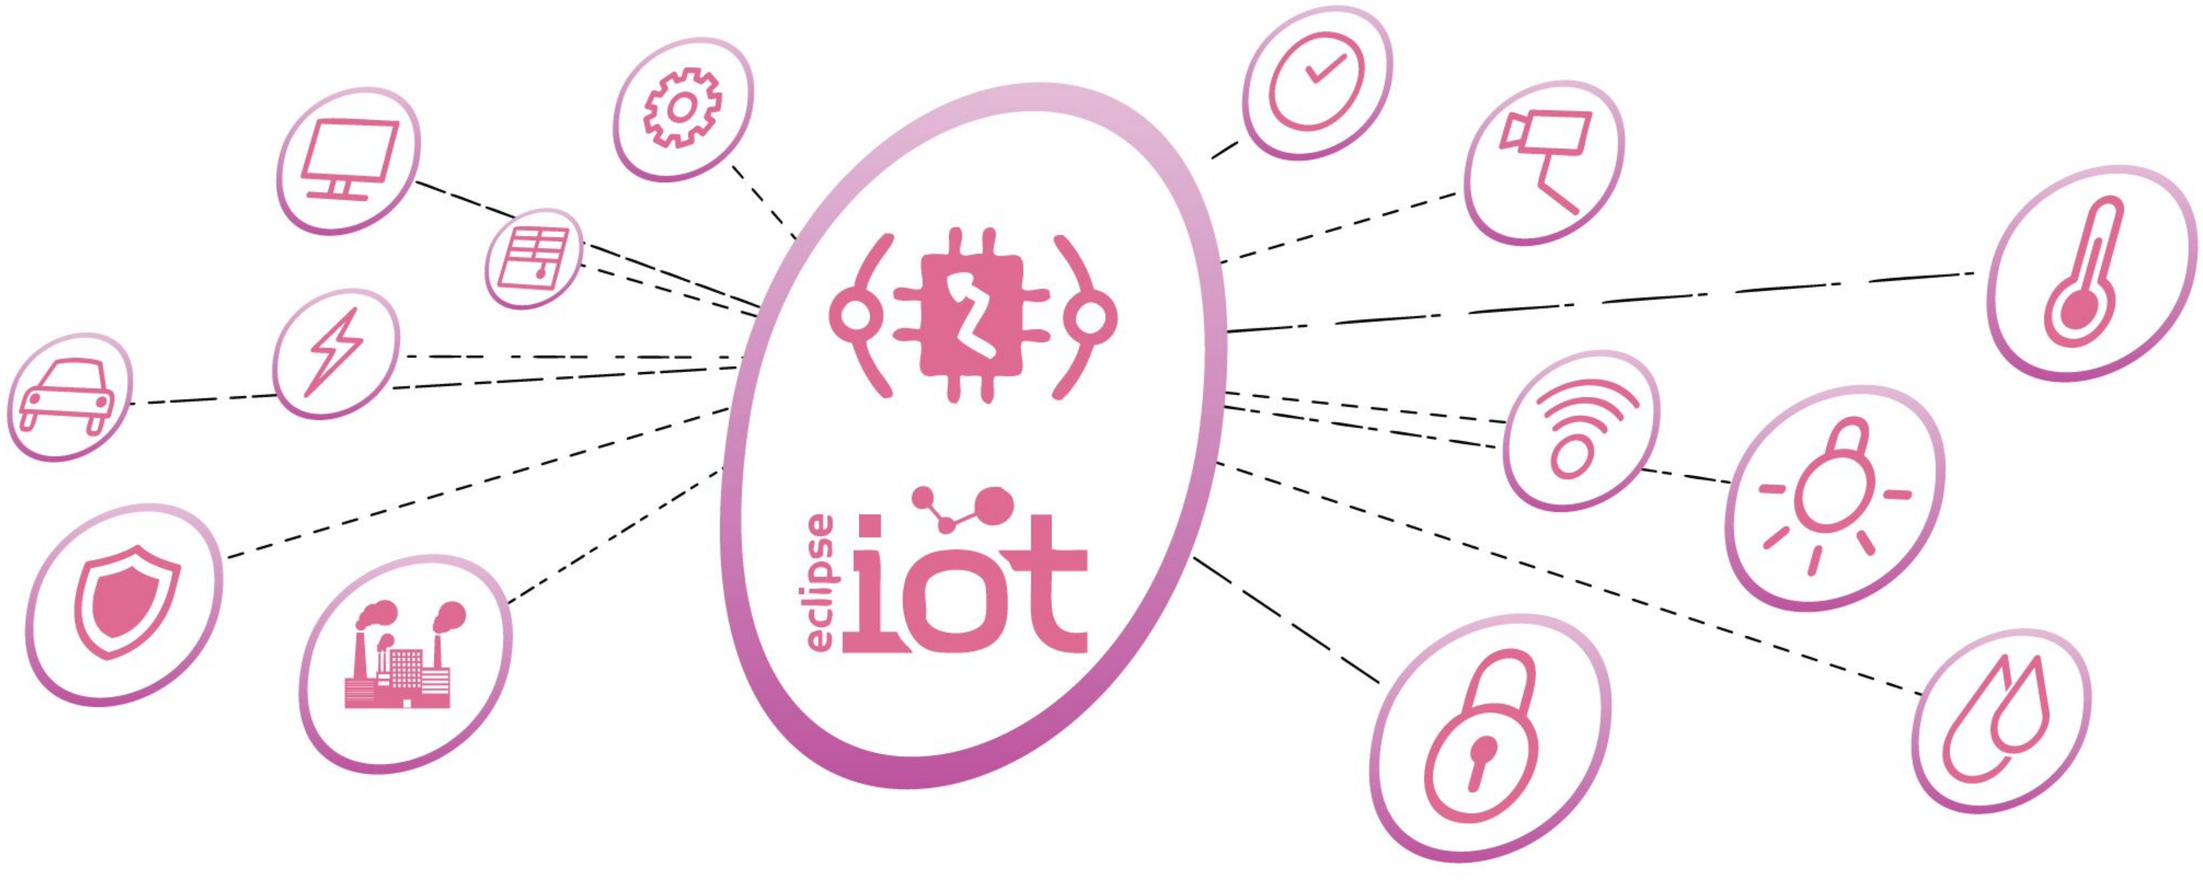
In the confluence of AI and quantum computing, neuromorphic systems stand as potent allies. Their brain-inspired designs resonate with AI's aspirations, and their adaptive architectures might dovetail with quantum's complexities. As IoT devices become ubiquitous, smarter, and more integrated into our lives, neuromorphic computing's relevance will only amplify, heralding a future where devices don't just connect but also cogitate.

Neuromorphic computing is more than just another tech buzzword; it presents a genuine paradigm shift. By drawing inspiration from the most complex structure known to us, the human brain, we're not merely aiming for better computing; we're redefining what it means to compute. END

By: Mir H.S. Quadri

The author is a research analyst with a specialisation in artificial intelligence and machine learning. He is the founder of Arkinfo, which focuses on the research and development of tech products using new age technologies. He shares a deep love for the analysis of technological trends and understanding their implications. Being a FOSS enthusiast, he has contributed to several open source projects.

How Open Source is Helping IoT Evolve



Let’s delve into the essence of IoT, explore its profound impact across diverse industries, and introduce the pivotal role of open source software in shaping its evolution.

IoT, short for the Internet of Things, encompasses a vast network of interconnected physical entities objects, devices, vehicles, and so on. These entities are equipped with sensors, software, and network connectivity, allowing them to autonomously collect, exchange, and analyse data. This enables them to interact with and respond to their environment. The fusion of the physical and digital realms redefines the way we gather, process, and harness information, ushering in a new era of connectivity and intelligence.

IoT’s significance in industries and open source software’s impact

IoT’s significance transcends mere technological innovation, promising to revolutionise a multitude of industries such as healthcare, agriculture, manufacturing, and transportation. In healthcare, IoT devices offer real-time patient monitoring, elevating the quality of care and enabling remote patient oversight. Agriculture benefits from IoT-powered sensors optimising crop management and sustainable irrigation practices, while manufacturing gets efficient through predictive maintenance and downtime reduction. These examples underscore IoT’s impact on precision, efficiency, cost-effectiveness, safety, and sustainability across sectors.

Meanwhile, open source software lies at the core of IoT’s rapid evolution. It adopts a collaborative ethos, providing freely accessible source code for software modification and distribution. Open source communities have been instrumental in driving IoT innovation, sharing knowledge, and creating versatile tools and platforms.

The choice of technologies and tools can significantly impact the success of a project. Open source technologies and tools have gained prominence for their flexibility, affordability, and vibrant communities.

Overview of key open source IoT platforms

Open source IoT platforms are the backbone of many IoT applications, providing the infrastructure for data collection, processing, and device management. Here are some popular open source IoT platforms.

OpenIoT

OpenIoT is a scalable and flexible open source middleware for the Internet of Things. It enables the integration of various IoT devices and data sources, making it a versatile choice for IoT applications.

```
# Example: Using OpenIoT for data
ingestion
from openiot import connect

# Connect to OpenIoT platform
iot = connect('mqtt://broker.example.
com')

# Publish IoT data
iot.publish('iot/sensors/temperature',
payload='25°C')
```

Eclipse IoT

Eclipse IoT is an ecosystem of open source projects that provide tools and frameworks for building IoT solutions. It includes projects like Eclipse Mosquitto (MQTT broker) and Eclipse Kura (IoT gateway).

```
// Example: Configuring Eclipse Mosquitto
as an MQTT broker
port 1883
listener 9001

# ...other configuration settings
```

ThingsBoard

ThingsBoard is an open source IoT platform that offers device management, data visualisation, and rule-based automation. It simplifies IoT application development and provides a rich set of widgets for creating dashboards.

```
// Example: Creating a dashboard widget
in ThingsBoard
const temperatureWidget = new
TbWidgetV2(
  'temperature-widget',
  {
    title: 'Temperature',
    datasource: self.ctx.datasource,
    dataKey: 'temperature',
    mainDataKey: 'temperature',
  }
);
```

Popular open source IoT protocols

Efficient communication protocols are crucial in IoT, and open source has provided several robust options. Here are two widely used open source IoT protocols.

MQTT (Message Queuing Telemetry Transport)

MQTT is a lightweight, open source protocol designed for efficient and reliable messaging in IoT. It follows a publish-subscribe model and is well-suited for constrained devices.

```
# Example: Using MQTT for IoT
communication
import paho.mqtt.client as mqtt

client = mqtt.Client()
client.connect("broker.example.com")

# Publish data
client.publish("iot/sensors/
temperature", "22.5°C")

# Subscribe to a topic
client.subscribe("iot/sensors/+/data")
```

CoAP (Constrained Application Protocol)

CoAP is an open source protocol designed for resource-constrained devices and low-power networks. It is similar to HTTP and is ideal for IoT applications that require a web-like interface.

```
// Example: Implementing CoAP server in
Java
CoapServer server = new CoapServer();
server.add(new
CoapResource("temperature") {
    @Override
    public void handleGET(CoapExchange
exchange) {
        // Respond with temperature data
        exchange.respond("22.5°C");
    }
});
server.start();
```

Examples of successful open source IoT projects

The open source community has birthed numerous IoT projects that exemplify the power of collaboration and innovation. Here, we'll dive into three successful open source IoT projects: Home Assistant, OpenHAB, and Node-RED. Let's explore how these projects leverage coding to create robust and user-friendly IoT solutions.

Home Assistant

Home Assistant is a highly popular and versatile open source home automation platform that has gained a massive following in the IoT community. It's designed to simplify the management and automation of various smart devices in your home. Here's why Home Assistant stands out.

Code example - Automation rule in Home Assistant (YAML)

```
# Example: Turn on the lights when
motion is detected
automation:
  - alias: "Turn on lights when motion
detected"
    trigger:
      platform: state
      entity_id: binary_sensor.motion_
sensor
      to: "on"
    action:
      service: light.turn_on
      entity_id: light.living_room
```

In this YAML code, an automation rule is defined to turn on the lights when motion is detected by the motion sensor. Home Assistant allows users to create such rules to automate various aspects of their smart home.

OpenHAB

OpenHAB is another prominent open

source home automation platform that emphasises flexibility and customisation. It provides a powerful environment for integrating and automating smart devices.

Code example - OpenHAB rules (DSL-based)

```
// Example: Turn on the heating when
the temperature drops below 20°C
rule "Temperature Control"
when
    Item temperature_sensor changed
then
    if (temperature_sensor.state < 20)
{
        heating_switch.sendCommand(ON)
    } else {
        heating_switch.sendCommand(OFF)
    }
end
```

OpenHAB uses a domain-specific language (DSL) for creating automation rules. In this example, the rule monitors the temperature sensor and controls the heating system accordingly.

Node-RED

Node-RED is a flow-based development tool for visual programming, ideal for creating IoT applications with a user-friendly approach. It simplifies the integration of devices, APIs, and online services. Here's what makes Node-RED stand out.

Code example - Node-RED flow for visual programming

```
[
    {
        "id": "1",
        "type": "inject",
        "payload": "{\"temperature\":
22.5}",
        "repeat": "10",
        "crontab": "",
        "once": false,
        "x": 150,
        "y": 220,
        "wires": [
            [
                "2"
            ]
        ]
    },
    {
        "id": "2",
        "type": "ui_gauge",
        "name": "Temperature",
        "group": "1",
        "order": 0,
        "width": 0,
        "height": 0,
        "gtype": "gage",
        "title": "Temperature",
        "label": "°C",
        "format": "{{value}}",
        "min": 0,
        "max": "40",
        "colors": [
            "#00b500",
            "#e6e600",
            "#ca3838"
        ],
        "x": 350,
        "y": 220,
        "wires": []
    },
    {
        "id": "3",
        "type": "ui_group",
        "name": "Sensor Data",
        "order": 1,
        "disp": true,
        "width": "6",
        "collapse": false
    },
    {
        "id": "4",
        "type": "ui_tab",
        "name": "Dashboard",
        "icon": "dashboard",
        "order": 1,
        "disabled": false,
        "hidden": false
    }
]
```

This shortened Node-RED flow simulates temperature data and displays it on a gauge in the Node-RED dashboard. It maintains the core functionality while reducing the code for brevity.

Benefits of open source in IoT

Open source is not just a development model; it's a philosophy that has significant advantages for IoT projects. In this section, we'll explore the key benefits of incorporating open source in the realm of IoT development.

Cost-effectiveness

- *Reduced development costs*

Building IoT solutions from the ground up can be a resource-intensive endeavour. Open source software eliminates the need to start from scratch, saving both time and money. Developers can leverage existing open source components and frameworks to accelerate development.

- *No licensing fees*

One of the most appealing aspects of open source is the absence of licensing fees. Traditional proprietary software often comes with substantial licensing costs. In contrast, open source allows organisations to use, modify, and distribute software freely, significantly reducing the financial burden on IoT projects.

Flexibility and customisation

- *Tailoring solutions to specific needs*

Open source empowers developers to tailor IoT solutions to meet specific requirements. Whether it's customising features, interfaces, or integrations, open source code can be modified to fit unique use cases, ensuring a precise fit for your project.

- *Avoiding vendor lock-in*

Proprietary solutions often result in vendor lock-in, making it challenging to migrate or adapt when needs change. Open source liberates organisations from such constraints, allowing for long-term flexibility and control over technology choices.

Community collaboration

- *Global developer community*

Open source thrives on community collaboration. A diverse and global community of developers contributes

to and enhances open source projects. This wealth of collective knowledge and expertise is a valuable resource for IoT projects, fostering innovation and problem-solving.

- *Rapid innovation*

The collaborative nature of open source accelerates innovation. Developers from different backgrounds and organisations work together to improve open source software continually. This results in faster development cycles, quicker adoption of new technologies, and the evolution of IoT solutions at a rapid pace.

Security and transparency

- *Code review and auditing*

Open source encourages peer review of code, which enhances security. When multiple eyes scrutinise code for vulnerabilities and bugs, it becomes more robust and resilient against potential threats. This continuous code review process bolsters the security of IoT solutions.

- *Mitigating vulnerabilities*

Security vulnerabilities can be quickly identified and addressed in open source projects. The transparency of open source code means that any flaws or weaknesses are more likely to come to light, allowing for prompt patching and mitigation. This reduces the risk of cyber threats compromising IoT systems.

Incorporating open source into IoT projects brings cost-effectiveness, flexibility, community collaboration, and enhanced security. These benefits empower organisations to create efficient, customisable, and secure IoT solutions that can evolve with changing needs. The synergy between open source and IoT continues to drive innovation and shape the future of connected technology.

Challenges and considerations

While open source offers numerous advantages in IoT development, it also presents unique challenges and considerations that need to be addressed effectively.

- ***Licensing and compliance issues***

Open source software comes with specific licensing requirements that must be adhered to. Failing to comply with these licences can lead to legal issues and reputational damage. It's crucial to have a clear understanding of the licensing terms associated with the open source components used in your IoT solution. Establish processes and policies to ensure compliance throughout the development and deployment phases. Regularly review licensing changes and updates to stay in line with legal requirements.

- ***Integration and interoperability challenges***

IoT environments often consist of a diverse range of devices, protocols, and technologies. Ensuring seamless integration and interoperability among these elements can be challenging. Open source solutions may need customisations to work cohesively within your IoT ecosystem. Carefully plan and test integrations, considering compatibility issues and data exchange protocols. Prioritise open standards to enhance interoperability and reduce complexity.

- ***Long-term support and maintenance***

IoT deployments are long-term endeavours, and open source components used in your solution require ongoing support and maintenance. The rapid pace of technological change can lead to dependencies on outdated or unsupported software. To address this challenge, establish a clear strategy for long-term support. Regularly update and patch open source components to address security vulnerabilities and ensure compatibility with evolving IoT standards. Consider investing in professional support or engaging with the open source community to maintain critical components.

- ***Balancing openness with security***

Open source's transparent nature is a double-edged sword when it comes to security. While it allows for code review and transparency, it also exposes vulnerabilities to potential threats. Striking a balance between openness and security is paramount. Implement robust security measures within your IoT solution, such as strong authentication, encryption, and access controls. Continuously monitor and audit open source components for security vulnerabilities. Engage with the open source community to report and address security issues promptly. Ensure that your development team is well-versed in security best practices to mitigate risks effectively.

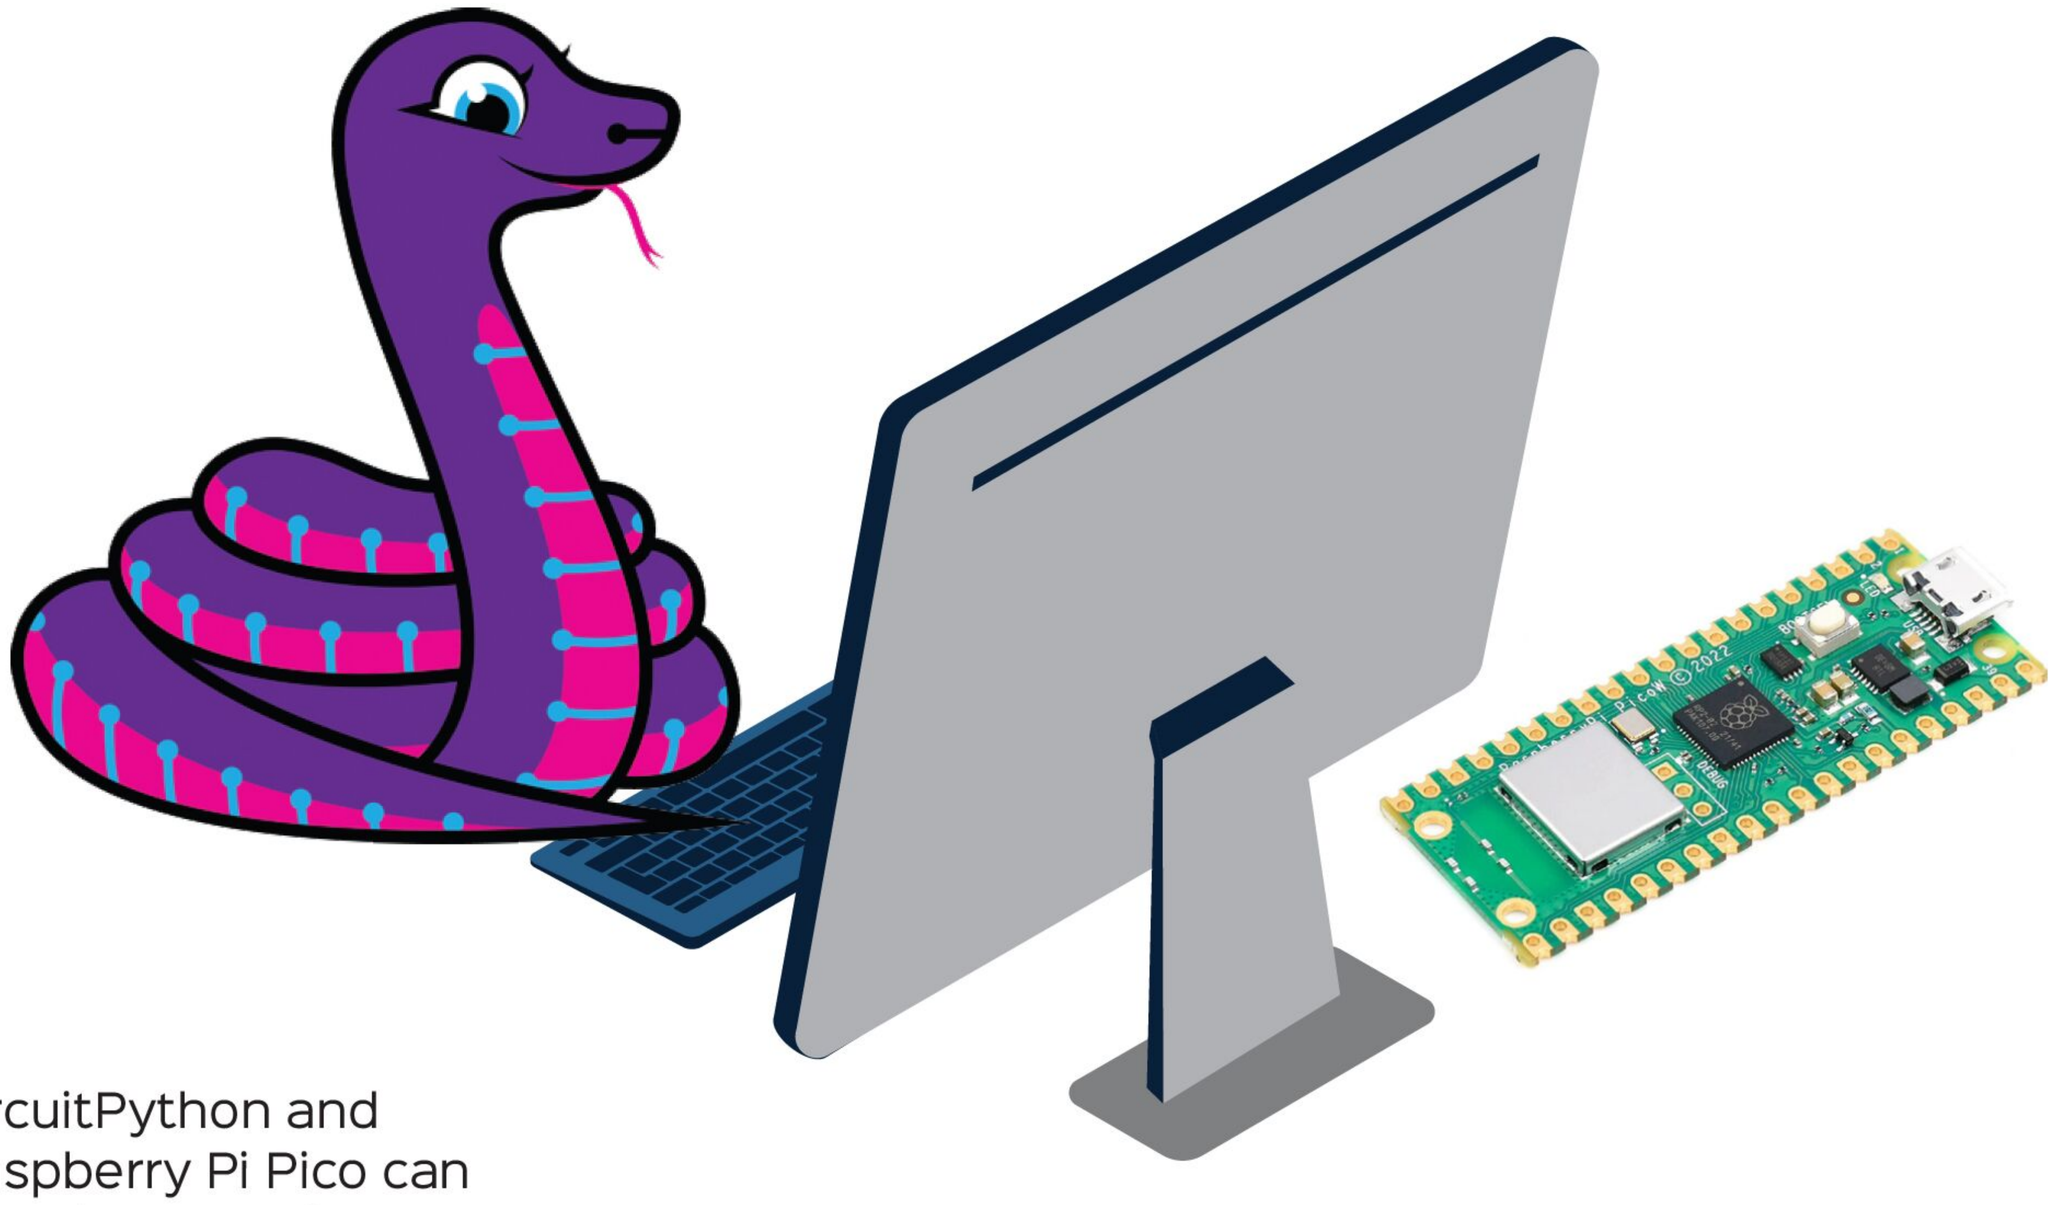
Precisely addressing these challenges and considerations is essential for harnessing the benefits of open source while mitigating potential risks in your IoT projects. A proactive approach to licensing, integration, long-term support, and security will contribute to the overall success and sustainability of your open source IoT solution.

Adopting open source empowers organisations to craft flexible, cost-effective, and secure IoT solutions. The future of IoT with open source appears promising, offering boundless opportunities for technological advancement and collaboration. It's time for organisations to wholeheartedly adopt open source principles in their IoT strategies, unlocking the full potential of the interconnected world. END

By: Dhaval Gajjar

The author is the CTO of Textdrip and the CEO of Pranshtech Solutions and WeTechnolabs Solutions. With a passion for technology and innovation, he explores the ever-evolving world of digital solutions, sharing insights and expertise to drive progress in the tech industry.

Leveraging the Power of CircuitPython for IoT Devices



CircuitPython and Raspberry Pi Pico can together transform everyday devices into smart wonders. This short tutorial will help you write a simple program for using CircuitPython and upload it to Raspberry Pi Pico.

The Internet of Things (IoT) refers to a network of devices capable of communicating with each other and the internet using unique identifiers. These devices have processing capabilities, high computing power, sensing units, and powerful embedded systems, allowing for applications as varied as ‘smart devices’, ‘smart cities’, and ‘smart schools’. IoT connects everyday objects like lights, monitors, cars, and doors to the internet, enabling the exchange of data and information among people, processes, and things. This data and information find applications in fields like healthcare, education, and agriculture. For instance, IoT can help monitor farming environments, prevent theft by tracking home activities, control machines in factories, and report power outages.

Some of the most popular IoT devices are Alexa and smartwatches, which are used for home automation and healthcare. These devices rely on hardware components called microcontroller units (MCUs) and microprocessor units (MPUs) to perform complex tasks. For example, Alexa uses the NXP i.MX RT106A MCU, enabling it to run the Alexa voice service directly on the device. It also incorporates the MediaTek MT8516 MPU, which supports voice assistants and machine learning applications.

Raspberry Pi Pico, the high-performance MCU

In this article, we will explore an MCU called Raspberry Pi Pico based on the RP2040 microcontroller chip by Raspberry Pi, and learn how to write a sample code for blinking an LED using CircuitPython. Raspberry Pi Pico is a low-cost, high-performance MCU board that uses the RP2040 chip. This chip boasts a dual-core ARM Cortex-M0+ processor, 264KB of SRAM, and 2MB of flash memory. Additionally, it offers 26 GPIO pins, three analogue inputs, and support for I2C, SPI, and UART protocols. Raspberry Pi Pico can run Python or C/C++ codes, making it remarkably easy to program and customise. Let me guide you through setting up your development environment, writing a simple firmware program, and uploading it to your Raspberry Pi Pico board.

Exploring Python for IoT

Python is a widely-used and versatile programming language with applications spanning various domains, including the Internet of Things (IoT). One of the devices that can be used for IoT projects is the Raspberry Pi Pico, as this low-cost microcontroller board can run Python code. To program the Pico with Python, we will use CircuitPython, a Python variant optimised for microcontrollers and embedded systems.

CircuitPython has many built-in libraries and modules that simplify the coding and interfacing with hardware components. To power our Pico device, we will need a power source that can provide 5V and 2A current source. Options include a battery pack or a power adapter with a micro-USB connector. For this article, I have used an Adafruit adapter, which is a USB wall charger compatible with the Pico board. You can also use any other power adapter that meets the specifications and has a micro-USB port.

Circuit connections

Here is a sample of a digital circuit that I created. It connects GPIO pin 22 to a resistor with a minimum resistance of 500 ohms. For a green LED circuit, you can use a 220-ohm resistor. This is because the frequency of red light is the highest in the visible light spectrum, and according to the following equation, energy is directly proportional to frequency:

$$Energy = h * frequency$$

This means that red LED bulbs heat up much faster than green LED bulbs. To prevent this, high-resistance resistors are needed. The other end of the resistor connects to the positive terminal of the LED bulb, marked by the LED's smaller leg. The negative end is connected to the GND (ground) pin of the Raspberry Pi Pico.

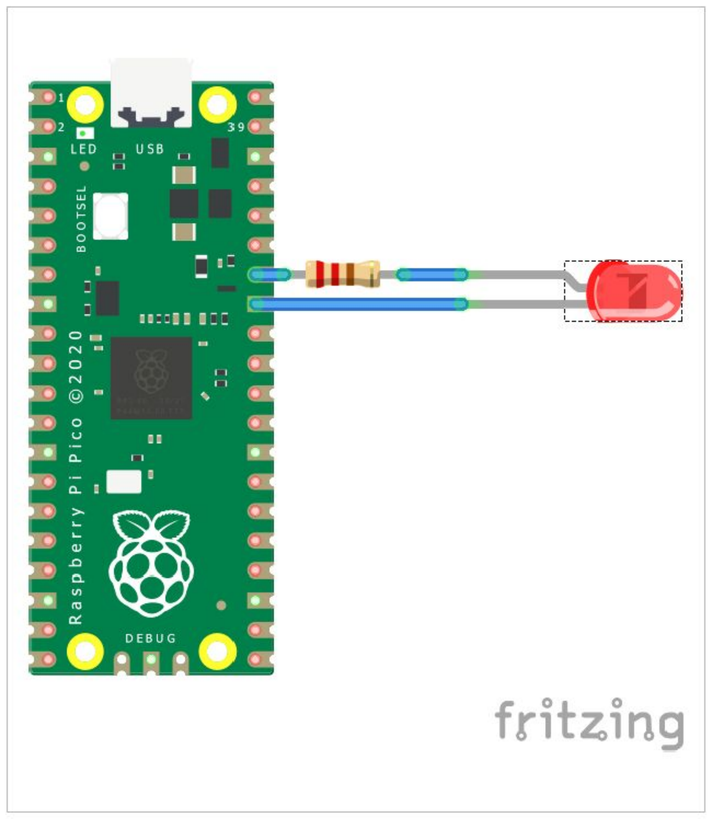


Figure 1: Raspberry Pi Pico with LED connected

Software configuration

To program the Raspberry Pi Pico with CircuitPython, follow these steps:

- Download the latest version of CircuitPython for Pico from the official website.
- Connect the Pico to the PC by holding down the BOOTSEL button while inserting the USB cable. This puts the Pico in bootloader mode, and a drive named RPI-RP2 will appear on the screen.
- Drag and drop the *adafruit_circuitpython_etc.uf2* file downloaded in step 1 to the RPI-RP2 drive. This installs CircuitPython on the Pico, and the drive disappears.
- A new drive called CIRCUITPY appears, which contains the files and folders that are required for CircuitPython to run on the board.

The main file we will use is *code.py*, which is the default CircuitPython code file executed when the board is powered on or reset. The *lib* folder contains the libraries or modules for our code.

Writing and running the code

To write and run our code, we can use any text editor or IDE supporting CircuitPython. In this case, we will use Thonny IDE, a simple and beginner-friendly IDE with a built-in CircuitPython interpreter. To use Thonny IDE, we need to select the CircuitPython interpreter from the *Tools* menu and choose Raspberry Pi Pico as the device.

In Thonny IDE, we can write our code in the editor window and save it as *code.py* on the CIRCUITPY drive. For example, we can write code to blink an LED light connected to pin 22 of the Pico board with a 0.5 second interval, as shown here:

```
Blink.py
1 import time
2 import board
3 import digitalio
4
5 led = digitalio.DigitalInOut(board.
GP22)
6 led.direction = digitalio.Direction.
OUTPUT
7
8 while True:
9    led.value = not led.value
10  time.sleep(0.5)
11  led.value = not led.value
12  time.sleep(0.5)
13
```

In this article, we have learned how to program the Raspberry Pi Pico with CircuitPython. We covered downloading and installing CircuitPython, using Thonny IDE to write and run code, and demonstrated how to blink an LED light connected to the Pico board.

CircuitPython is a powerful and user-friendly programming language that empowers us to create a wide range of IoT projects using the Pico board and other hardware components. END

By: Anisha Ghosh

The author is a cybersecurity researcher and an active contributor to open source communities and repositories. She is interested in developing novel, scalable and secure systems.

Choose the Right Cloud Platform for Implementing IoT Protocols

Working with IoT protocols like MQTT, AMQP, and CoAP on cloud platforms is essential for developing scalable and efficient IoT applications. The choice of the programming platform will depend on factors like project requirements, existing skills, and target devices. Leveraging the appropriate libraries and cloud services can enable seamless integration of IoT devices with cloud-based applications.



The Internet of Things (IoT) has ushered in a technological revolution, transforming the way we interact with the world around us. Its applications are vast and diverse, spanning numerous industries and sectors.

One of the most prevalent and impactful use cases of IoT can be found in the realm of smart home automation. IoT devices like smart thermostats, lighting systems, security cameras, and appliances enable homeowners to control and monitor their homes remotely, enhancing convenience, energy efficiency, and security. Some major IoT use cases have been catalogued in Table 1.

Industrial IoT (IIoT) is pivotal in the industrial sector, optimising operations through predictive maintenance, asset tracking, quality control, and supply chain management. This leads to reduced downtime, improved productivity, and data-driven decision-making, enhancing competitiveness.

In healthcare, IoT aids remote patient monitoring and wearable health devices, resulting in better health outcomes and reduced costs. Transportation and logistics benefit from IoT through vehicle tracking, fleet management, and traffic optimisation, boosting safety and efficiency. There are a range of

Table 1: Prominent and notable IoT use cases

Smart home automation • Smart thermostats • Smart lighting systems • Home security systems • Smart appliances	**Industrial IoT (IIoT)** • Predictive maintenance for machinery • Asset tracking and management • Quality control in manufacturing • Supply chain optimisation
Healthcare • Remote patient monitoring • Wearable health devices • Medication adherence monitoring • Smart prosthetics	**Transportation and logistics** • Vehicle tracking and fleet management • Traffic management and optimisation • Predictive maintenance for vehicles • Autonomous vehicles
Agriculture • Precision agriculture • Livestock monitoring • Soil and environmental monitoring • Automated irrigation systems	**Smart cities** • Smart traffic management • Waste management optimisation • Environmental monitoring • Smart street lighting
Retail • Inventory management and tracking • Customer behaviour analysis • Smart shelves and checkout systems • Personalised shopping experiences	**Energy management** • Smart grid technology • Energy consumption monitoring • Demand response systems • Renewable energy integration
Environmental monitoring • Wildlife tracking and conservation • Water quality monitoring • Forest fire detection and prevention • Pollution control	**Smart buildings** • Building automation and HVAC control • Energy-efficient lighting systems • Occupancy and space utilisation monitoring • Security and access control
Supply chain • Real-time tracking of goods in transit • Temperature and humidity monitoring for perishables • Inventory optimisation • Blockchain-based supply chain transparency	**Security and surveillance** • Smart cameras and video analytics • Intrusion detection systems • Fire and smoke detection • Access control systems
Smart farming • Crop and soil monitoring • Livestock tracking and health monitoring • Automated harvesting and sorting • Pest and disease management	**Health and fitness** • Fitness wearables • Smart scales and blood pressure monitors • Remote health consultations • Medication dispensers
Water management • Leak detection and prevention • Water quality monitoring • Irrigation control • Flood prediction and response	**Manufacturing** • Process automation and optimisation • Quality control and defect detection • Worker safety monitoring • Inventory tracking in real-time
Smart tourism • Location-based services and recommendations • Smart transportation options • Tourist information kiosks • Cultural heritage preservation and monitoring	**Education** • Smart classrooms and interactive learning • Student attendance monitoring • Campus security and access control • Asset tracking for educational institutions
Wearable technology • Smartwatches • Fitness trackers • Augmented reality (AR) glasses • Healthcare monitoring devices	**Sports and entertainment** • Sports performance analytics • Smart stadiums and fan engagement • Ticketing and access control • Equipment tracking and maintenance

applications in other domains too.

- **Agriculture:** Embraces IoT for precision farming, improving crop yields and sustainability
- **Smart cities**: Leverage IoT for traffic management, waste optimisation, and environmental monitoring, making urban areas more efficient and sustainable
- **Retail:** Uses IoT for inventory management and personalised shopping experiences
- **Energy management**: This is enhanced by IoT, with smart grids and energy consumption monitoring promoting sustainability
- **Environmental monitoring:** IoT sensors aid in tracking wildlife, water quality, and pollution levels, contributing to conservation efforts
- **Smart buildings:** Utilise IoT for energy optimisation and security
- **Supply chain management:** Relies on real-time tracking and inventory optimisation
- **Security and surveillance:** IoT enhances security and surveillance with advanced camera and sensor technologies

IoT also plays a crucial role in water management, tourism, education, wearables, and sports and entertainment, offering a wide range of benefits across these domains.

Key protocols in IoT implementations

In the world of IoT implementations, protocols play a pivotal role in enabling devices to communicate, exchange data, and function seamlessly within the IoT ecosystem. One of the foundational protocols is MQTT (Message Queuing Telemetry Transport), known for its lightweight and efficient nature. MQTT is ideal for scenarios where low bandwidth and minimal power consumption are crucial. This makes it a popular choice for IoT devices, especially in remote and resource-constrained environments.

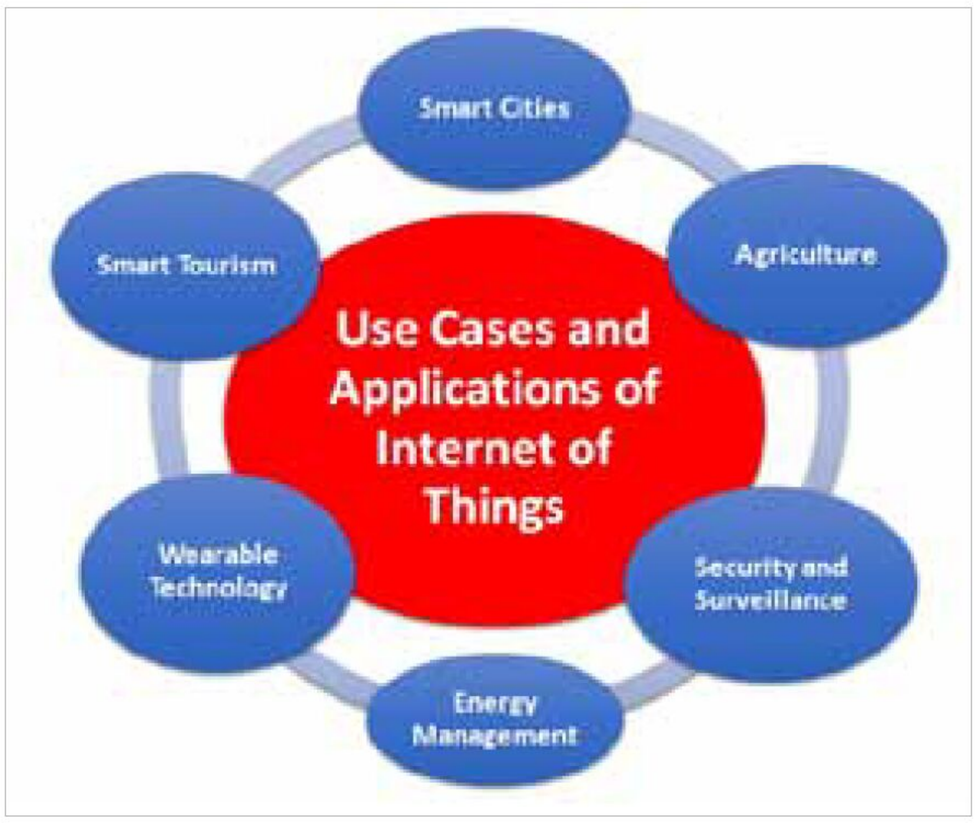


Figure 1: Use cases and applications of Internet of Things

HTTP (Hypertext Transfer Protocol), a protocol well-known in web communication, also plays a significant role in IoT. With the emergence of RESTful APIs, IoT devices can communicate with web servers over HTTP, enabling seamless integration with web applications and services. This simplifies data exchange and provides a familiar interface for developers.

On the other hand, CoAP (Constrained Application Protocol) is designed explicitly for IoT devices with limited resources. It operates like HTTP but with optimisations for low-power and constrained environments. CoAP is well-suited for IoT implementations where efficiency and scalability are paramount.

For IoT applications that require real-time communication and low latency, WebSocket is a critical protocol. It enables full-duplex communication between devices and servers, allowing for instantaneous data exchange. This protocol is particularly valuable in scenarios like remote control and monitoring.

AMQP (Advanced Message Queuing Protocol) is another protocol that finds its place in IoT implementations, especially in scenarios requiring reliable message queuing and routing. It provides a robust framework for message-oriented middleware, ensuring messages reach their destinations efficiently and safely.

In industrial and enterprise IoT applications, Modbus and BACnet are prevalent protocols. Modbus is used extensively in industrial control systems and automation, facilitating communication between devices like sensors and programmable logic controllers. BACnet, on the other hand, is focused on building automation and HVAC systems, ensuring interoperability and efficient data exchange in smart buildings.



Figure 2: Key protocols and implementation scenarios for IoT

Lastly, Bluetooth and Zigbee are wireless communication protocols that dominate IoT implementations in smart homes and personal area networks. Bluetooth is well-known for its versatility and widespread adoption in consumer devices, while Zigbee excels in low-power, mesh network applications, often found in home automation setups.

In IoT, the choice of protocol depends on various factors such as the nature of the application, resource constraints, scalability requirements, and interoperability needs. The seamless integration of these key protocols is essential for creating a robust and interconnected IoT ecosystem that enables devices to communicate and deliver the benefits of IoT to users and businesses alike.

Working with IoT protocols on cloud

Working with IoT protocols such as MQTT, AMQP, and CoAP on cloud platforms is crucial for building scalable and efficient IoT applications. These protocols facilitate communication between devices and cloud services, allowing for data exchange, control, and management of IoT devices. Developers can work with these protocols using various programming platforms like Python, Node.js, Java, and more, depending on their preferences and project requirements.

Python is a versatile and widely-used language for IoT development. To work with IoT protocols in Python, developers can use libraries and frameworks such as Paho MQTT for MQTT, the RabbitMQ library for AMQP, and CoAPthon for CoAP. These libraries enable developers to establish connections, publish and subscribe to topics, and handle messages efficiently. When combined with cloud services like AWS IoT, Azure IoT Hub, or Google Cloud IoT, Python allows seamless integration of IoT devices with cloud platforms.

Node.js, known for its asynchronous and event-driven architecture, is also a popular choice for IoT development. Developers can use libraries like MQTT.js for MQTT, AMQP10 for AMQP, and CoAP for CoAP. Node.js's non-blocking I/O model makes it well-suited for handling numerous IoT device connections simultaneously, making it a great choice for IoT gateway applications that connect devices to the cloud.

Java is a robust and platform-independent language for building IoT applications. It offers libraries like Eclipse Paho for MQTT, Apache Qpid for AMQP, and Californium for CoAP. Java's strong community support and extensive libraries make it a reliable choice for developing IoT applications that need to run on different platforms and devices.

Other programming platforms such as C/C++ can be used for IoT protocol implementations when working with resource-constrained devices. Libraries like Eclipse Mosquitto for MQTT, RabbitMQ for AMQP, and microcoap for CoAP enable efficient communication with cloud platforms.

Working with MQTT, AMQP and CoAP

Working with cloud-based MQTT, AMQP, and CoAP protocols using Python is a common approach for developing IoT applications that require efficient communication between devices and cloud services. Below, we'll outline how to work with each of these protocols in a cloud environment using Python.

MQTT (Message Queuing Telemetry Transport): MQTT is a lightweight and widely-used protocol for IoT communication. To work with MQTT in a cloud-based environment using Python, you can follow these steps.

1. Install the Paho MQTT library, a popular MQTT client library for Python, using pip:

```
pip install paho-mqtt
```

Figure 3 presents a cloud based MQTT broker that serves as a critical component in an IoT project.



Figure 3: Cloud based MQTT for IoT scenarios

However, the process of setting up, troubleshooting, monitoring, and sustaining such a broker can be intricate and time-intensive. MQTTHQ is an intricately designed, load-balanced, multi-node MQTT broker cluster. It has been meticulously crafted to offer a steadfast and exceptionally reliable broker environment for the development of IoT products and applications. MQTTHQ extends its support to both TCP and WebSocket connections. The project integration can instantaneously transmit and receive messages via the public.mqtthq.com broker, all without the need to delve into coding intricacies.

2. Import the Paho MQTT library in your Python script:

```
import paho.mqtt.client as mqtt
```

The Paho MQTT library stands as a cornerstone in the world of IoT and messaging protocols. This open source library, developed by the Eclipse Foundation, empowers developers to effortlessly integrate MQTT communication into their applications.

Offering support for multiple programming languages, including Python, Java, and JavaScript, Paho simplifies the process of connecting devices, exchanging messages, and managing MQTT connections. Its lightweight and efficient design makes it a go-to choice for IoT projects, ensuring seamless communication between devices and MQTT brokers. With Paho MQTT, developers can harness the power of MQTT for building robust and scalable IoT solutions with ease.

3. Create an MQTT client instance and set up the connection to the MQTT broker on the cloud platform:

```
client = mqtt.Client(client_id="your_
client_id")
client.connect("mqtt_broker_address",
port=1883)
```

Use the client to publish and subscribe to MQTT topics. You can send data from your IoT devices to the cloud by publishing to a specific topic and subscribe to topics to receive commands or updates from the cloud.

AMQP (Advanced Message Queuing Protocol): AMQP is a robust and feature-rich messaging protocol. To work with AMQP in a cloud-based environment using Python:

1. Install an AMQP library like `pika` for Python using pip:

```
pip install pika
```

2. Import the `pika` library in your Python script:

```
import pika
```

3. Create a connection to the AMQP broker on the cloud platform:

```
connection = pika.
BlockingConnection(pika.
ConnectionParameters('amqp_broker_
address'))
channel = connection.channel()
```

3. Declare queues and exchanges for message routing, ensuring they match the setup on your cloud-based AMQP service.
4. Use the 'channel' object to publish and consume messages from the queues. You can publish data from your IoT devices and consume messages sent from the cloud.

CoAP (Constrained Application Protocol): CoAP is a lightweight and efficient protocol designed for resource-constrained IoT devices. To work with CoAP in a cloud-based environment using Python:

1. Install a CoAP library like 'aiocoap' for Python using pip:

```
pip install aiocoap
```

2. Import the `aiocoap` library in your Python script:

```
import aiocoap
```

3. Create a CoAP client instance and set up the connection to the CoAP server (usually provided by your cloud-based service):

```
context = await aiocoap.Context.create_
client_context()
uri = "coap://coap_server_address/
resource"
```

4. Use the client to send CoAP requests (GET, PUT, POST, etc) to retrieve or update data on the cloud server. CoAP follows a RESTful architecture, so you interact with resources using standard HTTP-like methods.

Working with MQTT, AMQP, and CoAP in a cloud-based environment using Python allows you to create powerful and scalable IoT applications. Ensure that you have the necessary credentials and access rights to connect to the cloud-based services, and that you properly configure your Python script to match the cloud platform's settings and requirements.

As technology advances and connectivity expands, IoT's influence will only grow, opening new horizons and opportunities for businesses and individuals alike.

When working with cloud platforms for IoT, consider using IoT-specific cloud services like AWS IoT Core, Azure IoT Hub, or Google Cloud IoT Core. These services simplify device registration, message routing, and device management, making it easier to integrate IoT devices with cloud applications. END

By: Dr Gaurav Kumar

The author is associated with various academic and research institutes for delivering expert lectures and conducting technical workshops on the latest technologies and tools.

Simplifying the World of Distributed Ledger Technology

We take a look at what managed blockchain is, its advantages, and its role in shaping the future of distributed ledger technology. We also explore key AWS managed blockchain services and their features, and find out how blockchain can be used to prevent land grabbing.



In the quiet state of Salamia, nestled amidst rolling hills and lush forests, a new chapter was about to unfold. The residents of the state had long lived in the shadow of a previous regime that had left behind a legacy of misdeeds and mistrust. However, hope blossomed as a new chief minister, Naresh Pamnani, assumed power with a determination to rectify past wrongs.

Naresh was a visionary leader, armed with a keen intellect and a heart fuelled by justice. From the moment he took office, he had a clear goal in mind: to tackle the pervasive issue of land grabbing that had plagued Salamia for years. Rumours had circulated about the lands being unfairly taken from rightful owners, leaving families displaced and communities torn apart.

With a team of technologists, legal experts, and community leaders by his side, Naresh introduced an innovative solution to the age-old problem – a blockchain-based land registry system. This system aimed to ensure the accuracy, transparency, and immutability of land ownership records.

Initially the people of Salamia were sceptical. But Naresh was relentless in his pursuit of justice. He launched a grassroots campaign, visiting villages, holding town hall meetings, and listening to the stories of those who had suffered due to land grabs. His empathy and determination resonated with the people, and slowly, their scepticism began to wane.

The blockchain-based land registry system was rolled out with great care and meticulous planning. Every land parcel in Salamia was digitised and represented as a unique digital token on the blockchain. The state's residents could access this decentralised ledger and verify ownership records themselves, eliminating the need for intermediaries and reducing the risk of fraud.

Smart contracts were introduced to automate land transactions. These contracts would hold funds in escrow

until all conditions were met, ensuring that transactions were fair, secure, and free from manipulation. Community members actively participated in the process, ensuring that the system truly served their needs.

As the months went by, the impact of Naresh's initiative became evident. Families were able to reclaim their ancestral lands, long lost to unscrupulous practices. Communities that had once been divided found themselves united by a shared hope for justice. The once-sceptical residents of the state now saw Naresh as a beacon of change, a leader who had turned their hopes into reality.

News of Salamia's success story spread far and wide, inspiring neighbouring states and even cities to consider similar blockchain-based solutions for land management. Blockchain had woven a tapestry of trust, enabling a new era of prosperity and unity.

Blockchain technology has been making waves in various industries, promising transparency, security, and efficiency in record-keeping and transactions. However, implementing and managing a blockchain network can be a complex and resource-intensive endeavour. To address these challenges, managed blockchain services have emerged, offering businesses a simplified and efficient way to harness the power of blockchain technology.

Managed blockchain

Managed blockchain is a cloud-based service that provides businesses with a fully managed infrastructure to deploy, operate, and scale blockchain networks without the need for extensive in-house expertise. These services are offered by major cloud providers such as Amazon Web Services (AWS), Microsoft Azure, and IBM Cloud, among others.

Key components of managed blockchain

Network management: Managed blockchain services provide a user-friendly interface for creating and managing blockchain networks. Users can set up their networks, define permissions, and configure nodes without delving into the intricacies of blockchain technology.

Node management: Users can easily add or remove nodes from their blockchain network, ensuring scalability as their needs evolve. This eliminates the need for manual node maintenance and reduces downtime.

Security and identity management: Managed blockchain platforms integrate robust security features, including identity and access management, to protect sensitive data and transactions. This reduces the risk of unauthorised access and fraud.

Chaincode (smart contracts) deployment: Users can deploy smart contracts on their blockchain networks effortlessly. These self-executing contracts automate complex processes, improving efficiency and reducing the potential for human error.

Monitoring and analytics: Managed blockchain services offer real-time monitoring and analytics tools, allowing users to gain insights into network performance, transaction history, and overall health.

Advantages of managed blockchain

Reduced complexity: One of the primary advantages of managed blockchain is its ability to abstract the complexity of blockchain technology. This makes it accessible to a broader range of businesses, even those without extensive IT resources or blockchain expertise.

Cost-efficiency: These services typically operate on a pay-as-you-go model, reducing upfront infrastructure costs. Businesses can scale resources up or down based on demand, optimising their spending.

Rapid deployment: Setting up a blockchain network through a managed service is significantly faster than building one from scratch. This allows businesses to quickly experiment with blockchain applications and respond to market demands promptly.

Scalability: Managed blockchain services offer seamless scalability, ensuring that businesses can adapt to increased transaction volumes and network growth without disruption.

Robust security: Leading cloud providers invest heavily in security measures, making managed blockchain networks highly secure. This reduces the risk of data breaches and unauthorised access.

Compliance and governance: These blockchain platforms often include tools for setting up governance models and ensuring compliance with industry-specific regulations.

Needless to say, managed blockchain has found applications in various industries, including supply chain management, financial services, healthcare, retail, energy, etc.

Amazon Web Services (AWS) offers a robust suite of managed blockchain services that allow organisations to easily create, deploy, and manage blockchain networks without the complexities typically associated with setting up and maintaining such networks. These services are designed to help businesses harness the power of blockchain technology for various use cases, including supply chain management, financial services, healthcare, and more.

Let's explore the key AWS managed blockchain services and their features.

Amazon Managed Blockchain

Amazon Managed Blockchain (AMB) is one of the core blockchain services offered by AWS. It allows users to create and manage scalable blockchain networks using popular open source frameworks like Hyperledger Fabric and Ethereum. AMB takes care of the underlying infrastructure and offers features such as network setup, member management, security, and monitoring.

***Ease of setup*:** Amazon Managed Blockchain simplifies the process of creating a blockchain network, allowing users to set up a network with just a few clicks through the AWS management console.

Network management: Users can create and manage multiple blockchain networks, each with its own set of members and policies.

Security: The service integrates with AWS identity and access management (IAM) to control access to resources securely. It also allows for encryption of data in transit and at rest.

Scalability: Amazon Managed Blockchain can automatically scale to accommodate changes in network activity and transaction volume, ensuring high availability.

Monitoring and analytics: The service provides real-time monitoring and insights into the performance of the blockchain network, helping users identify and address issues quickly.

Smart contracts: It supports the deployment of smart contracts (Chaincode for Hyperledger Fabric and Solidity for Ethereum) to automate business processes.

Integration: Amazon Managed Blockchain can be integrated with other AWS services like Amazon S3, AWS Lambda, and Amazon CloudWatch for building comprehensive blockchain applications.

Amazon Quantum Ledger Database (Amazon QLDB)

While not a traditional blockchain, Amazon QLDB is a fully managed ledger database service designed to provide a transparent and immutable journal for tracking changes to data over time. Key features of Amazon QLDB include:

Immutable ledger: It maintains an immutable and tamper-evident record of all changes made to data, making it suitable for applications that require a transparent and auditable data history.

Data verification: The service provides cryptographic verification of data integrity, ensuring that data remains unchanged and can be trusted.

SQL-like queries: Users can query the ledger using a familiar SQL-like language, simplifying data retrieval and analysis.

Managed service: QLDB is fully managed by AWS, reducing operational overhead, and making it easy to scale as needed.

Integration: It can be integrated with other AWS services and applications to create comprehensive data-driven solutions.

AWS blockchain templates

AWS provides templates for popular blockchain frameworks, such as Ethereum and Hyperledger Fabric, to help users quickly deploy and configure blockchain networks on AWS infrastructure. These templates simplify the setup process.

Amazon Managed Blockchain for Ethereum

This service is specifically designed for Ethereum-based blockchain networks. It offers Ethereum-compatible APIs, smart contract support, and simplified management of Ethereum nodes.

Amazon Managed Blockchain for Hyperledger Fabric

This service is tailored for Hyperledger Fabric-based blockchain networks. It provides tools for creating, configuring, and managing Hyperledger Fabric nodes and networks.

Amazon Aurora with Amazon QLDB integration

AWS offers integration between Amazon Aurora (a relational database service) and Amazon QLDB, allowing businesses to create hybrid applications that combine the benefits of traditional databases with the immutability of ledger technology.

Ethereum on AWS

Ethereum's versatility has led to its widespread adoption for various use cases, including decentralised finance (DeFi), non-fungible tokens (NFTs), supply chain management, identity verification, and more. However, as with any technology, it's important to understand its capabilities and limitations before building or investing in projects on the Ethereum blockchain.

Key features of the Ethereum blockchain

Ethereum is a decentralised, open source blockchain platform that enables the creation and execution of smart contracts and decentralised applications (DApps). It was proposed in late 2013 by Vitalik Buterin and development began in early 2014, with the mainnet (live blockchain network) launched on July 30, 2015. The key features of the Ethereum blockchain include:

Smart contracts: Ethereum introduced the concept of smart contracts, which are self-executing contracts with the terms of the agreement directly written into code. These contracts automatically execute when specific conditions are met, without the need for intermediaries.

Decentralised applications (DApps): Ethereum allows developers to build decentralised applications on its platform. These DApps can serve a wide range of purposes, from financial services to gaming and more. They are typically open source and operate on the blockchain, providing transparency and security.

Ether (ETH): Ether is the native cryptocurrency of the Ethereum network. It is used to pay for transaction fees and computational services on the platform. It's also often seen as a digital asset and investment.

Decentralised autonomous organisations (DAOs): Ethereum supports the creation of DAOs, which are organisations run by code and the consensus of their members. They can

make decisions and manage funds autonomously, based on the rules defined in their smart contracts.

Consensus mechanism: Ethereum currently operates on a proof of stake (PoS) consensus mechanism known as Ethereum 2.0 (Eth2). This is an upgrade from the original proof of work (PoW) mechanism, which is still used in Ethereum 1.0. Eth2 aims to improve scalability, security, and energy efficiency.

EVM and Solidity: Ethereum introduced the Ethereum Virtual Machine (EVM), a runtime environment that enables smart contracts to be executed. Solidity is the most commonly used programming language for writing smart contracts on Ethereum.

Interoperability: Ethereum has supported the development of various standards and protocols, such as ERC-20 (fungible tokens) and ERC-721 (non-fungible tokens), which have contributed to the growth of the blockchain ecosystem.

Upgrades: Ethereum is continually evolving. Major upgrades are being rolled out to address scalability, security, and sustainability challenges.

It's important to note that Ethereum, like other blockchains, requires participants (nodes) to validate transactions and maintain the network. This decentralised nature helps ensure security, censorship resistance, and immutability of the blockchain's data.

In order to work with Ethereum on AWS, we may consider using Amazon Managed Ethereum (AME) or deploying our Ethereum nodes manually on AWS services such as EC2 instances or Kubernetes clusters.

Amazon Managed Ethereum (AME)

AME is a fully managed Ethereum service provided by AWS. It is designed to simplify the deployment and management of Ethereum infrastructure.

It allows the creation of Ethereum networks quickly without the need to manage the underlying infrastructure. We can use AME to deploy Ethereum nodes, interact with smart contracts, and build decentralised applications (DApps) on the Ethereum network.

Solution for land grabbing prevention

Blockchain technology, when used in conjunction with legal and societal reforms, can help mitigate land grabbing by increasing transparency, improving property record management, and enhancing trust. Let us see how Ethereum blockchain technology can contribute to the prevention of land grabbing.

Immutable land records: Ethereum's blockchain offers immutability, meaning that once data is recorded on the blockchain, it is extremely difficult to alter or delete it. Land records stored on Ethereum can provide a tamper-resistant history of ownership and transactions, reducing the risk of fraudulent changes to land titles.

Transparent land registries: Using Ethereum-based smart contracts, governments or organisations can create transparent land registries that are accessible to the public.

Identity verification: Implementing a blockchain-based identity verification system can ensure that only legitimate property owners or authorised parties can initiate land transactions. This can reduce the chances of unauthorised land sales.

Smart contracts for transactions: Ethereum's smart contracts can be used to automate and validate land transactions. When a transaction is initiated, the smart contract can verify the legitimacy of the buyer and seller, and ensure that all legal requirements are met before transferring ownership.

Timestamping and notarisation: Blockchain can be used for timestamping and notarisation of land-related documents, creating a provable history of documents related to land transactions.

Community involvement: Blockchain can facilitate community involvement and consensus in land management decisions. Community members can collectively validate and approve land transactions to prevent unilateral decisions that may lead to land grabbing.

Legal and regulatory frameworks: While blockchain can provide technical solutions, it must be integrated into a supportive legal and regulatory framework. Governments and authorities should recognise blockchain-based land records and transactions as legally binding.

It's important to recognise that blockchain technology is not a silver bullet for preventing land grabbing. It should be part of a broader strategy that includes legal reforms, community engagement, and good governance practices. Additionally, blockchain systems should be designed with careful consideration of privacy, security, and scalability to be effective in land management. END

Acknowledgements

The author would like to thank Balakrishnan Srinivasan (CTO - hybrid cloud application modernisation services, IBM India) for providing encouragement to write this article, Gopalakrishna U. (associate partner, IBM India) for granting the opportunity to gain real-time experience on this subject, and Harsh Mehta (delivery manager, IBM India) for the dedicated efforts in releasing this article for publishing.

By: Venkateswaran N.

The author is a lead architect, containerisation and modernisation, AWS practice, IBM India.

Latest Trends in DevOps and What's Next

Enterprises must embrace DevOps and its evolutionary trends to ensure they stay ahead of the competition and remain customer-centric. Here are some insights into the pivotal DevOps trends of automation, security, AI integration, and more, and how they benefit organisations.



In recent years, DevOps has become a crucial component of software development and delivery, enabling teams to collaborate and deliver high-quality applications at a faster pace. As organisations increasingly adopt DevOps practices, the industry is witnessing new trends and advancements that are shaping its future. One significant trend is the shift towards automation, enabled by tools like Kubernetes and Jenkins. The adoption of artificial intelligence and machine learning is also gaining momentum, as organisations seek to optimise their DevOps processes further.

Another trend is the rise of observability, which enables teams to gain valuable insights into their application's performance and user behaviour. This approach involves the collection and analysis of various data points, such as logs, metrics, and traces, to identify potential issues and improve overall system health.

In the coming years, DevOps is expected to continue evolving, with a greater emphasis on security and compliance. The use of DevSecOps practices, which integrate security measures into the development and delivery process, is becoming increasingly crucial as cyber threats become more sophisticated.

Importance of keeping up with DevOps trends

Stay ahead of the competition: Staying up-to-date with the latest trends in DevOps can help organisations stay ahead of the competition and stay relevant in the market.

Improve productivity: DevOps trends often focus on improving automation and efficiency in the development process, which can result in improved productivity for your team.

Enhanced collaboration: Keeping up with DevOps trends can also lead to better collaboration between different teams and departments within an organisation.

Better quality software: Following DevOps trends can also help organisations deliver better quality software, through improved testing processes and faster release cycles.

Reduce costs: DevOps trends often emphasise reducing the time and costs associated with development and deployment, which can lead to significant cost savings for your organisation.

Current DevOps trends

DevOps is an essential methodology that helps organisations to improve their software development life cycle by aligning their development and operational teams. The DevOps culture is continuously evolving, and with more companies adopting this approach to improve their software delivery, we can see some significant DevOps trends emerging in the technology world. Let's take a look at some of the exciting and innovative trends that are driving DevOps forward.

Automation: Automating repetitive tasks such as server configuration, application deployment, testing, and monitoring is becoming increasingly popular to minimise human errors and speed up the delivery process.

Continuous delivery: Continuous delivery is aimed at ensuring that the software is always in a deployable state, and any changes made to the codebase can be released to production immediately.

Microservices: Rather than monolithic architecture, microservices use a collection of loosely coupled services to provide flexible, scalable, and independent development processes.

Cloud-native: Cloud-native emphasises deploying applications in the cloud instead of relying on traditional on-premises infrastructure.

Security-as-Code: Security is baked into the CI/CD pipeline so that automated testing detects vulnerabilities before they become threats.

DevSecOps: DevSecOps integrates security practices into the DevOps pipeline, so security becomes everyone's responsibility, rather than of a separate team.

Infrastructure as Code (IaC): Infrastructure as Code is about defining and managing infrastructure using code to ensure reproducibility, consistency, and scalability.

GitOps: GitOps is a new approach to continuous delivery that uses Git as the single source of truth for managing infrastructure and applications.

Observability: Observability is the ability to monitor, understand, and troubleshoot complex systems in real-time.

Serverless computing: Serverless computing is an event-driven architecture that allows developers to deploy code without worrying about infrastructure management.

Automation and Infrastructure as Code (IaC)

Automation and Infrastructure as Code (IaC) are two closely related concepts that have revolutionised the way businesses build and manage their IT infrastructure. Automation refers to the use of tools and technologies to automate various IT operations, such as provisioning, configuration, testing, and deployment.

With automation, businesses can reduce the time and effort required to perform routine tasks, improve the accuracy and reliability of their IT operations, and free up IT teams to focus on more strategic initiatives. By automating various processes, businesses can also reduce costs and improve efficiency, enabling them to deliver products and services faster and more reliably.

Infrastructure as Code, on the other hand, refers to the practice of using code to manage and provision infrastructure. With IaC, businesses can use software code to automate the deployment and configuration of their infrastructure, making it more scalable, repeatable, and reliable. IaC allows businesses to define their infrastructure as code, which can be easily versioned, tested, and reviewed, enabling them to make changes quickly and efficiently.

IaC has several benefits, including reducing the risk of human error by automating repetitive tasks, making it easier to manage complex infrastructure, and improving collaboration and communication between different teams. Additionally, IaC can help businesses achieve regulatory compliance, as they can easily track and audit changes to their infrastructure.

Security integration into DevOps

DevOps focuses on collaboration and communication between different departments such as software development, testing, and operations to ensure the seamless and continuous delivery of software products.

However, security has often been an afterthought in the DevOps process, with many organisations prioritising speed and agility over security. This approach can result in security vulnerabilities in the software products released, leading to data breaches, cyber-attacks, and loss of user trust, which can significantly impact an organisation's reputation and financial standing.

To address this challenge, security integration into DevOps is becoming a critical component of software development processes. This approach involves integrating security practices, tools, and workflows into the DevOps cycle to ensure that security is at the forefront of the software development process.

One of the key benefits of security integration into DevOps is that it helps to identify potential security risks early during the development process. This early detection enables security teams to work closely with the development team to mitigate any potential vulnerabilities before software products are released.

Another benefit is that security is accessible to all departments involved in the software development process, and not just the security team. This approach reduces the siloed nature of security and promotes a culture of shared responsibility for security across the organisation.

To implement security integration into DevOps, organisations can adopt various practices such as:

- Implementing security testing and vulnerability assessments into the CI/CD pipeline to identify and mitigate security vulnerabilities early in the development process
- Providing security-focused training to all stakeholders involved in the software development process
- Implementing security automation and orchestration tools to streamline security workflows and improve collaboration between teams
- Incorporating security requirements into the software development process as part of the requirements gathering phase.

Shift towards cloud-native architecture

With businesses relying on cloud computing more than ever before, cloud-native architecture has become a key focus for companies looking to deliver a more agile and scalable IT infrastructure.

Cloud-native architecture focuses on using cloud computing technologies to create complex, distributed systems. This approach emphasises the use of containers and microservices, which break down applications into smaller, more manageable parts, making them easier to deploy and maintain.

The shift towards cloud-native architecture is driven by several factors. First, businesses are looking to modernise their applications and IT infrastructure, and the cloud provides the ideal platform for this. By moving to the cloud, businesses can take advantage of the scalability and flexibility of cloud computing technologies, enabling them to deliver better services to customers and stakeholders.

Second, there is a growing need for speed and agility in software development. With cloud-native architecture, developers can work more quickly and efficiently, deploying changes to applications faster and with less disruption. This gives businesses a competitive edge, enabling them to respond quickly to changing market conditions and customer needs.

Third, moving to cloud-native architecture can help businesses reduce their IT costs, as they no longer need to maintain their own infrastructure and can instead rely on cloud providers to manage it for them. This allows organisations to redirect their resources to other areas of the business.

Adoption of microservices

Microservices are a type of software architecture where applications are broken down into smaller, independent services that work together to deliver a larger functionality.

One of the key benefits of microservices is the ability to scale and update individual services without affecting the entire system. This allows for easy and less costly updates and upgrades, and can reduce downtime. Additionally, microservices can increase resilience by allowing services to fail independently, rather than causing an entire application to fail.

To adopt microservices, businesses typically start by breaking down their monolithic application into smaller services. This can be a time-consuming process, but it allows for a more modular architecture that can adapt to changing requirements.

Another important component of microservices is the use of APIs to communicate between services. This allows for greater flexibility and interoperability, making it easier to integrate with other systems and applications.

To support microservices, businesses may need to adopt new tools, such as containerisation and orchestration platforms like Kubernetes or Docker. These technologies enable the deployment and management of microservices at scale, making it easier to manage a complex system of services.

AI and machine learning integration

One of the primary benefits of integrating AI and machine learning into DevOps is improved speed and efficiency. AI and machine learning algorithms can automate repetitive and time-consuming tasks, allowing developers to focus on more complex aspects of software development. Additionally, these algorithms can analyse large amounts of data, such as performance metrics and user behaviour data, and identify areas where improvements can be made. This enables organisations to make data-driven decisions and optimise their DevOps processes.

Another benefit of integrating AI and machine learning into DevOps is improved quality and reliability. By automating testing processes, AI and machine learning algorithms can quickly identify bugs and other issues before they become major problems. This ensures that software is developed and released with as few errors as possible, reducing the need for costly bug fixes and improving customer satisfaction.

Furthermore, AI and machine learning algorithms can help organisations to detect and respond to security threats quickly and accurately.

They can monitor and analyse network traffic, detect anomalies, and flag suspicious behaviour, allowing organisations to take quick and effective action to mitigate risks.

As AI and machine learning continue to evolve and improve, we can expect to see even greater advancements in DevOps automation and optimisation.

Increased focus on observability, monitoring, and logging

In recent years, there has been an increased focus on observability, monitoring, and logging in the DevOps community. As organisations adopt modern practices for delivering software, it is important to have visibility into the performance and behaviour of applications and infrastructure.

Observability is the ability to measure the internal state of a system based on its outputs. This concept has gained traction in the DevOps world because it helps teams understand the function of software systems in real-time. By tracking metrics such as response time, error rates, and user satisfaction, teams can quickly identify and diagnose issues.

Monitoring is the process of collecting data about systems and applications in order to detect, diagnose, and resolve issues. This is especially important in complex environments where multiple services and components work together. By monitoring logs, metrics, and events, DevOps teams can identify potential problems and take proactive measures to prevent them.

Logging is the practice of recording system events and activity for later analysis. Logs are crucial for understanding the context in which issues occur, debugging complex problems, and improving overall system performance. With the rise of microservices and distributed systems, logging has become an essential component of observability.

The increased focus on observability, monitoring, and logging in DevOps is driven by several factors. First, the rise of cloud computing and containerisation has made it easier to deploy and scale applications, but it also introduces complexity. Second, DevOps teams are under pressure to deliver high-quality software at a faster pace, which requires a more efficient approach to monitoring and diagnosing issues. Finally, regulatory requirements and security concerns have made it essential to have visibility into how systems are performing and who is accessing them.

In a nutshell, observability, monitoring, and logging are critical elements of modern DevOps practices. By implementing these practices, organisations can improve their ability to detect and diagnose issues, improve system performance, and ensure compliance with regulatory requirements.

Implications of current trends

The current trends in DevOps have significant implications for businesses.

Automation: Automation tools are being used for various activities, including testing, deployment, and configuration management. Automation leads to better quality products, faster product development, and reduced time-to-market.

Cloud computing: The cloud offers scalability, flexibility, and cost savings to businesses. Cloud computing also enables teams to work remotely and collaborate effectively.

Microservices: Microservices architecture offers scalability, resilience, and flexibility to businesses.

Continuous integration and continuous deployment (CI/CD): CI/CD helps businesses to release new features and fixes rapidly, ensure quality through automated testing, and reduces the risk of production failures.

Security: The adoption of DevSecOps practices is necessary to ensure secure product development and deployment.

What's next in DevOps?

Enterprises are keen to embrace the next-generation DevOps practices to enhance the agility, efficiency, and scalability of their software development, management, and deployment processes. Here are some of the significant developments that can shape the future of DevOps.

Securing DevOps: Security is increasingly becoming an essential concern in the DevOps arena. DevOps teams need to understand the vulnerabilities and ensure the security of the application's infrastructure and code throughout the development and deployment cycle. Implementing security measures, such as automation, collaboration, and continuous testing, can improve the security of the DevOps pipeline.

Artificial intelligence and machine learning: AI and ML can automate repetitive tasks, detect and prevent defects, and enhance decision-making capabilities. The integration of AI and ML in the DevOps workflow can enable teams to identify areas for improvement and optimise the development and deployment process continually.

Cloud-native DevOps: Cloud-native DevOps leverages the cloud infrastructure's scalability, agility, and flexibility to enable rapid software development and deployment. This approach involves building applications using cloud-native technologies, such as containers and microservices, and integrating them with the DevOps pipeline for maximum productivity.

Agile DevOps: Agile DevOps combines agile methodologies with DevOps practices, allowing teams to deliver software rapidly, frequently, and with high quality. This approach involves collaboration, continuous delivery, and embracing change, making DevOps more agile while maintaining its core principles.

The road ahead

For a digital business, keeping up with DevOps trends is of utmost importance. DevOps practices are rapidly evolving and changing, and staying ahead of these changes can help organisations gain a competitive edge and stay relevant in their industry.

Here are some reasons why it's crucial to keep up with DevOps trends.

Improved efficiency: The latest DevOps trends are geared towards improving your team's efficiency and productivity. With tools like automation, continuous delivery, and cloud-based infrastructure, you can streamline your development processes, reduce wait times, and deliver products to market faster.

Better collaboration: DevOps trends encourage better collaboration and communication between teams. When developers, operations staff, and other stakeholders work closely together, it promotes shared ownership and a sense of camaraderie, which can translate into better code quality, faster troubleshooting, and more efficient problem-solving.

Enhanced security: Security is a top concern for any digital business, and DevOps trends are focusing more than ever on ensuring that security is built into the development and deployment process. By staying up to date with the latest trends in security, you can ensure that your applications are less vulnerable to cyber threats and that they meet industry standards and compliance regulations.

Increased innovation: DevOps gives organisations the ability to experiment more and iterate faster to bring new products or features to the market. Keeping up with DevOps trends allows you to stay on top of new technologies and methodologies that can potentially unlock new ways of doing things and drive innovation.

In conclusion, DevOps trends are constantly advancing, and keeping up with them can give you a strategic advantage in the market. By embracing these trends and incorporating them into your operations, you can improve efficiency, enhance security, drive innovation, and ultimately deliver better value to customers. END

By: Dr S. Balakrishnan

The author heads the Department of Computer Science and Business Systems at Sri Krishna College of Engineering and Technology, Coimbatore, Tamil Nadu.

OSFY Magazine Attractions During 2023-24	
Month	**Theme**
March 2023	Security, Network Management and Monitoring
April 2023	Open Source Programming (Languages and tools)
May 2023	Cloud Special: Everything from management to implementation
June 2023	AI, Deep learning and Machine Learning
July 2023	Database management and Optimisation
August 2023	Mobile/Web App Development, Optimisation and Security
September 2023	DevOps Special
October 2023	Blockchain and Open Source
November 2023	Open Source and IoT and Edge
December 2023	All About Data Management
January 2024	Containers and Managing Containers
February 2024	Open Source on Windows and Best in the world of Open Source (Tools and Services)

Revolutionising Land Management with Blockchain Technology

Blockchain promises a transparent, secure, and efficient land management system, fostering trust and empowering citizens.

Blockchain is a digital, distributed and decentralised network to store information in a tamper-proof form, providing an automated way to enforce trust among different participants. An open distributed ledger can record all transactions between different parties efficiently and in a verifiable and permanent way. It captures, manipulates, and builds consensus among participants in the network.

Each block is uniquely connected to the previous blocks via a digital signature, which means that making a change to a record without disturbing the previous records in the chain is not possible, thus rendering the information tamper-proof.

Here, I propose a traceable method for registering personal information on the blockchain, giving individuals control over their information and ensuring it is only shared with authorised parties. This solution has the potential to improve privacy and security and give individuals the ability to track and monitor their personal information.

What governments around the world are doing

According to *Forbes*, the Republic of Georgia has declared that they would use blockchain technology to validate property-related government transactions. Countries like Sweden and the Honduras are also developing similar blockchain-based systems, for enabling secured e-governance.

Recently, Dubai announced that they would put all their records pertaining to land registry on the blockchain. Dubai Land Department (DLD) claims to be the first government department anywhere in the world to adopt blockchain for such high-level tasks.

According to the *Times of India*, 76 land grab cases were filed in the courts in the country since 2018 and 50% of these were in 2022 alone.

As per the *SCC Online* blog, the Supreme Court of India has observed that if a state government is so conscious or interested in taking action against land grabbers, it is the same government's discretion to bring or

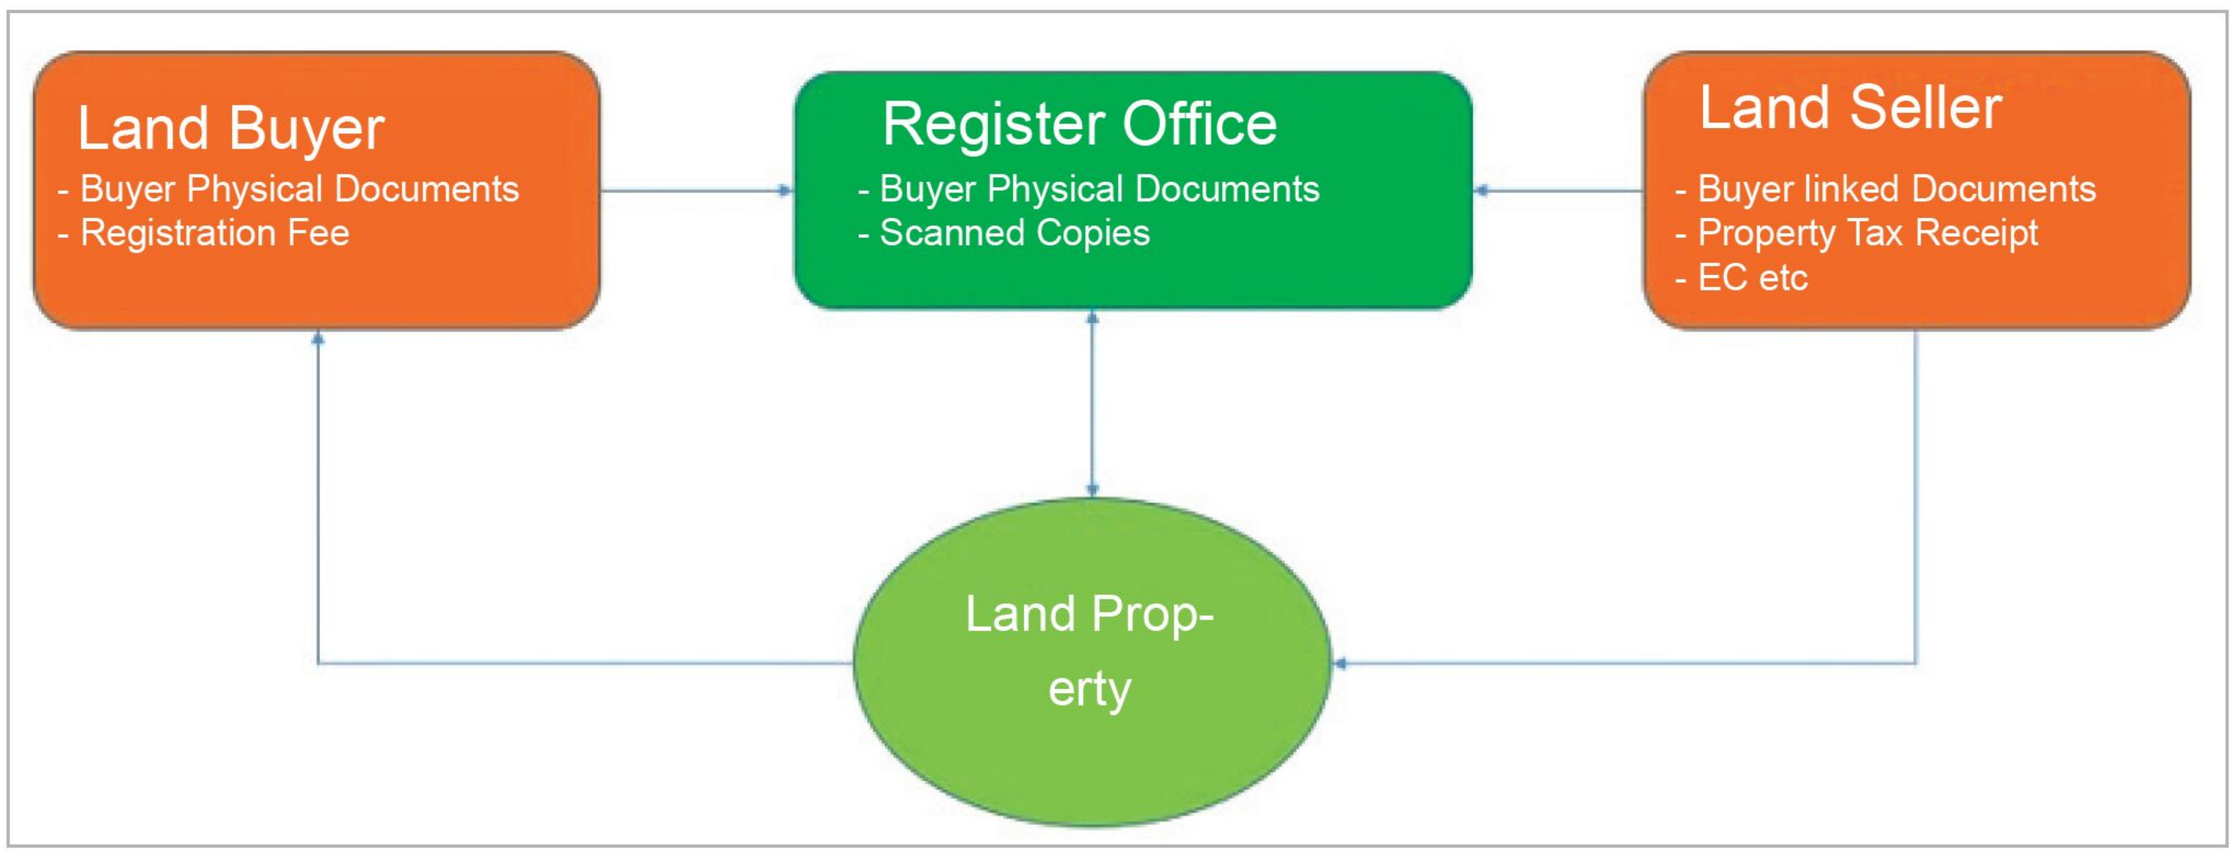


Figure 1: Current land management process

better the appropriate legislation with the definition of 'land grabber', 'land grabbing' and 'land grabbing cases'.

Andhra Pradesh has become the first state in India to pilot blockchain technology in two departments, and plans to deploy it across the administration.

Current state of the land management system

Land management entails the registration and management of all land deeds undertaken in any geographical location of a governing authority. Various government departments and local bodies maintain these records. The registration of land records does not authenticate the 'land title', which sometimes leads to title disputes and litigation.

The current land registration system in India involves a paper-based approach that is susceptible to fraud, errors, and security risks. The process is also time-consuming and involves huge documentation, which contributes to fraudulent activities.

One survey estimated that 70% of landowners are vulnerable to manipulation by individuals or mediators with stakes, who challenge the validity of their ownership. There are a huge number of examples across India where land has been grabbed by forging records.

Of late, there is an increase in complaints of irregularities and corruption in updating land records. Many Indians are eagerly waiting for corrections in their land records.

In today's scenario, the seller and buyer meet at the registration office, share the physical documents in the presence of the Registrar, and complete the land registration. The transaction primarily happens through the exchange and verification of physical copies.

The major issues with land registration in India are:

- Unclear land titles, leading to huge GDP loss
- Land grabbing
- Duplicate land records
- Land records are maintained in silos; hence, updating and verifying the records is a challenge
- Poor record-keeping, with stakeholders often dealing with inaccurate data
- The database method used for storing data is inadequate in terms of data security
- Multiple land web portals, and none of them provide a single source of truth

All the above facts demand the government to develop a secured, tamper-proof, paperless, and long-lasting online solution to manage land registration and land records. The best solution for this is implementation of blockchain technology-based land management.

Blockchain-based land management solution for the future

Blockchain technology helps establish trust, transparency and accuracy in maintaining land records and building a robust land management system. Implementation of blockchain in land records provides:

- Integrity
- Transparency
- Elimination of intermediaries
- Speed
- Elimination of documentation work
- Minimisation of fraud

The blockchain-enabled land management system is a decentralised public ledger spread over a distributed network that records every transaction associated with a land property in an autonomous and efficient manner. It enhances the accuracy and security of land records.

Blockchain works by validating transactions through a distributed network to create a permanent, verified, and unalterable ledger of information. Figure 2 shows the end-to-end flow of the blockchain-based land registration process.

The steps followed in this process are:

- The seller requests that the invoice be published by the purchaser. Seller stores the relevant factoring agreement in a smart contract on a public blockchain.
- Purchaser examines the smart contract, notices that the invoice has been factored and pays the invoice amount.
- After obtaining the signature of the purchaser and seller in the sale deed, the scanned document is moved into the blockchain network to create a block.
- The certificates issued by the revenue department are stored in the blockchain and can be used by other agencies like the bank for any verification process during a land transaction.
- Verification of the details is done using the blockchain data. After the approval of transactions in the respective database, such as completion of registration or approval of a loan by the bank, the transaction details are stored in the blockchain.

Once the block is created, it cannot be tampered with or edited. The chain of blocks is created every time the property title is changed from one person to another.

Also, the existing history of transactions on a piece of land first needs to be inserted into the blockchain after approval by revenue functionaries in the state. The approved data can be digitally signed and stored.

Blockchain technology helps to set up:

Tamper-proof land records: Land records are secure, tamper-proof and immutable. This reduces the scope of forgery or manipulation.

Unified registration system: A single digital platform removes bureaucratic silos and helps in driving greater efficiency and reducing inaccuracies/forgeries.

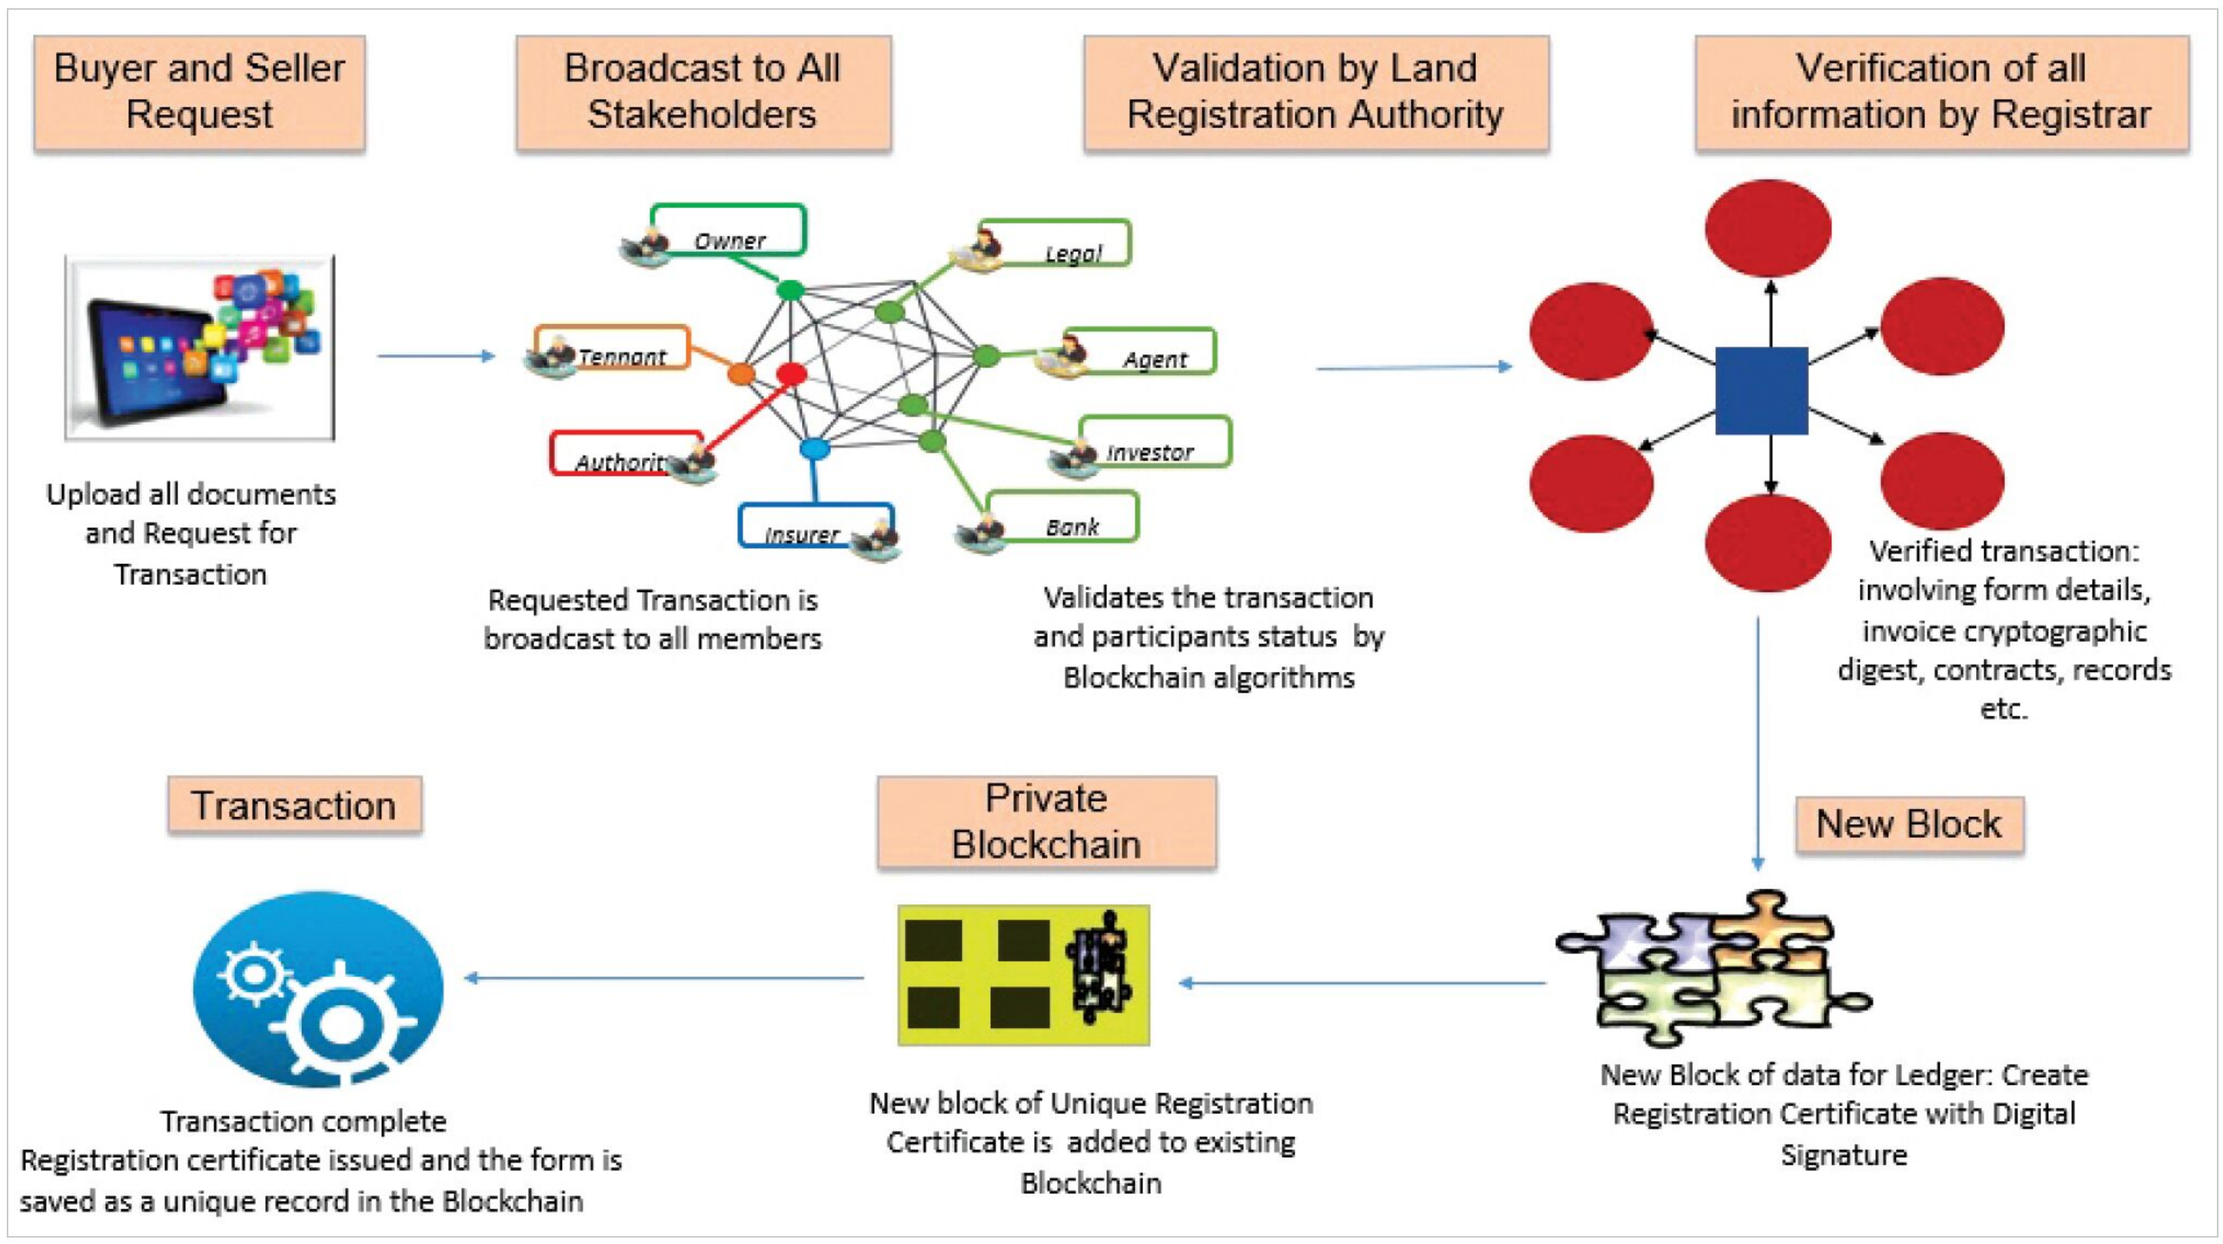


Figure 2: Future land management process

Transparent registration process: Automating the process of registration using smart contracts does away with mediators who are vulnerable to corrupt practices.

Benefits of blockchain technology in land management system

One of the major benefits of using blockchain is the reduction in the cost of storing physical records of any land property. Other benefits are:

- Greater transparency, as transaction histories become more transparent using blockchain technology
- No need for trusted authorities like notaries to provide attested copies of documents
- Citizens are assured that their land ownership cannot be changed by spurious persons
- Enhanced security – blockchain is far more secure than other record-keeping systems
- Elimination of error-handling through real-time tracking of transactions with no double spending
- Improved traceability
- Trusted recordkeeping
- Purchaser need not depend on any unreliable personnel/agency to verify the authenticity of the document provided by the seller
- Increased efficiency and speed, reducing settlement time to mere seconds by removing intermediaries
- Reduced cost and complexity due to the elimination of expensive proprietary infrastructure
- Full automation of transactional processes, from payment through settlement

Blockchain-based government applications prevent illegal activities involved in land transactions. All contracts and ownership details are maintained in a decentralised way, eliminating the need for physical intervention. Data transactions can be tracked easily, which improves the overall security of the system.

Blockchain is a unique and universal technology that helps streamline and automate nearly all citizen services and legal contracts, while increasing the transparency and effectiveness of government services. END

Acknowledgements

The author would like to thank Santosh Shinde of BTIS, Enterprise Architecture division of HCL Technologies Ltd for giving the required time and support in many ways in bringing this article as part of architecture practice efforts.

By: Dr Gopala Krishna Behara

The author is an enterprise architect in the BTIS Enterprise Architecture division of HCL Technologies Ltd. He has a total of 27 years of experience in the IT industry.

Disclaimer: The views expressed in this article are that of the author and HCL does not subscribe to the substance, veracity, or truthfulness of the said opinion.

R Series: Caret Package

In this 25th article in the 'R, Statistics and Machine Learning' series, we explore the caret package, which is a framework for building machine learning models in R.



The caret package provides miscellaneous functions for training classification and regression models. It allows you to tune parameters, use resampling techniques and make predictions from the models.

We will use R version 4.2.2 installed on Parabola GNU/Linux-libre (x86-64) for the R code snippets.

```
$ R --version

R version 4.2.2 (2022-10-31) -- "Innocent and Trusting"
Copyright (C) 2022 The R Foundation for Statistical Computing
Platform: x86_64-pc-linux-gnu (64-bit)

R is free software and comes with ABSOLUTELY NO WARRANTY.
You are welcome to redistribute it under the terms of the
GNU General Public License versions 2 or 3.

For more information about these matters, see https://www.
gnu.org/licenses/.
```

Package installation

You can install and load the caret package using the following commands:

```
> install.packages("caret")
Installing package into '/home/shakthi/R/x86_64-pc-linux-gnu-
library/4.1'
…

 * building package indices
 * installing vignettes
 * testing if installed package can be loaded from temporary
location
 * checking absolute paths in shared objects and dynamic
libraries
 * testing if installed package can be loaded from final
location
 * testing if installed package keeps a record of temporary
installation path
 DONE (caret)

> library(caret)
Loading required package: ggplot2
Loading required package: lattice
```

We will install and use a few additional packages to demonstrate the use of the *caret* package. The *earth* package implements multivariate adaptive regression splines (MARS). You can install and load this library using the following commands:

```
> install.packages("earth")
Installing package into '/home/shakthi/R/x86_64-pc-linux-gnu-
library/4.1'
...

 * building package indices
 * testing if installed package can be loaded from temporary
location
 * checking absolute paths in shared objects and dynamic
libraries
 * testing if installed package can be loaded from final
location
 * testing if installed package keeps a record of temporary
installation path
 DONE (earth)

> library(earth)
Loading required package: Formula
Loading required package: plotmo
Loading required package: plotrix
Loading required package: TeachingDemos
```

The *mda* package provides mixture and flexible discriminant analysis, MARS, and vector-response smoothing splines. You can install and load the *mda* library using the following commands:

```
> install.packages("mda")

Installing package into '/home/shakthi/R/x86_64-pc-linux-gnu-
library/4.1'
...
 ** checking absolute paths in shared objects and dynamic
libraries
 * testing if installed package can be loaded from final
location
 * testing if installed package keeps a record of temporary
installation path
 DONE (mda)

> library(mda)
Loading required package: class
Loaded mda 0.5-4
```

Glass data set

The *mlbench* package contains numerous machine learning benchmark problems and data sets from the UCI repository. Let us use the Glass data set that contains 214 observations of chemical analysis of seven different types of glass.

```
> library(mlbench)
> data(Glass)
```

The *set.seed()* function is used to reset the random number generation. You can then take a sample data from the Glass data set for training and test validation as shown below:

```
> set.seed(103)
> train <- sample(1:dim(Glass)[1], 150)

> data <- Glass[ train, ]
> test <- Glass[-train, ]
```

The structure of the Glass data is given below:

```
> str(data)
'data.frame':  150 obs. of  10 variables:
 $ RI  : num  1.52 1.52 1.52 1.52 1.52 ...
 $ Na  : num  13.6 13.7 13.1 13.4 13.7 ...
 $ Mg  : num  3.87 3.84 3.45 3.49 0 3.5 3.26 3.61 0 3.48 ...
 $ Al  : num  1.27 0.72 1.76 1.52 0.56 1.48 2.22 1.54 3.04
1.35 ...
 $ Si  : num  72 71.8 72.5 72.7 74.5 ...
 $ K   : num  0.54 0.17 0.6 0.67 0 0.6 1.46 0.66 6.21 0.64
...
 $ Ca  : num  8.32 9.74 8.38 8.08 10.99 ...
 $ Ba  : num  0 0 0 0 0 0 1.63 0 0 0 ...
 $ Fe  : num  0.32 0 0.17 0.1 0 0 0 0 0 0 ...
 $ Type: Factor w/ 6 levels "1","2","3","5",..: 2 1 3 2 2 2 6
2 4 1 ...
```

The *densityplot()* of the training data set is shown in Figure 1.

bagFDA()

The *bagFDA()* function is a wrapper function that uses multivariate adaptive regression splines (MARS) to implement flexible discriminant analysis (FDA). It accepts the following arguments:

Argument	Description
x	A data frame or matrix of examples
y	A matrix or data frame of outcomes
weights	Weights for examples (Default: 1)
B	Number of bootstrap samples
keepX	Boolean to retain original training data
formula	Represented as y ~ x1 + x2 + …
data	Data frame from which variables are fetched
subset	An index vector listing the samples to be used
na.action	A function to be applied if data has NA

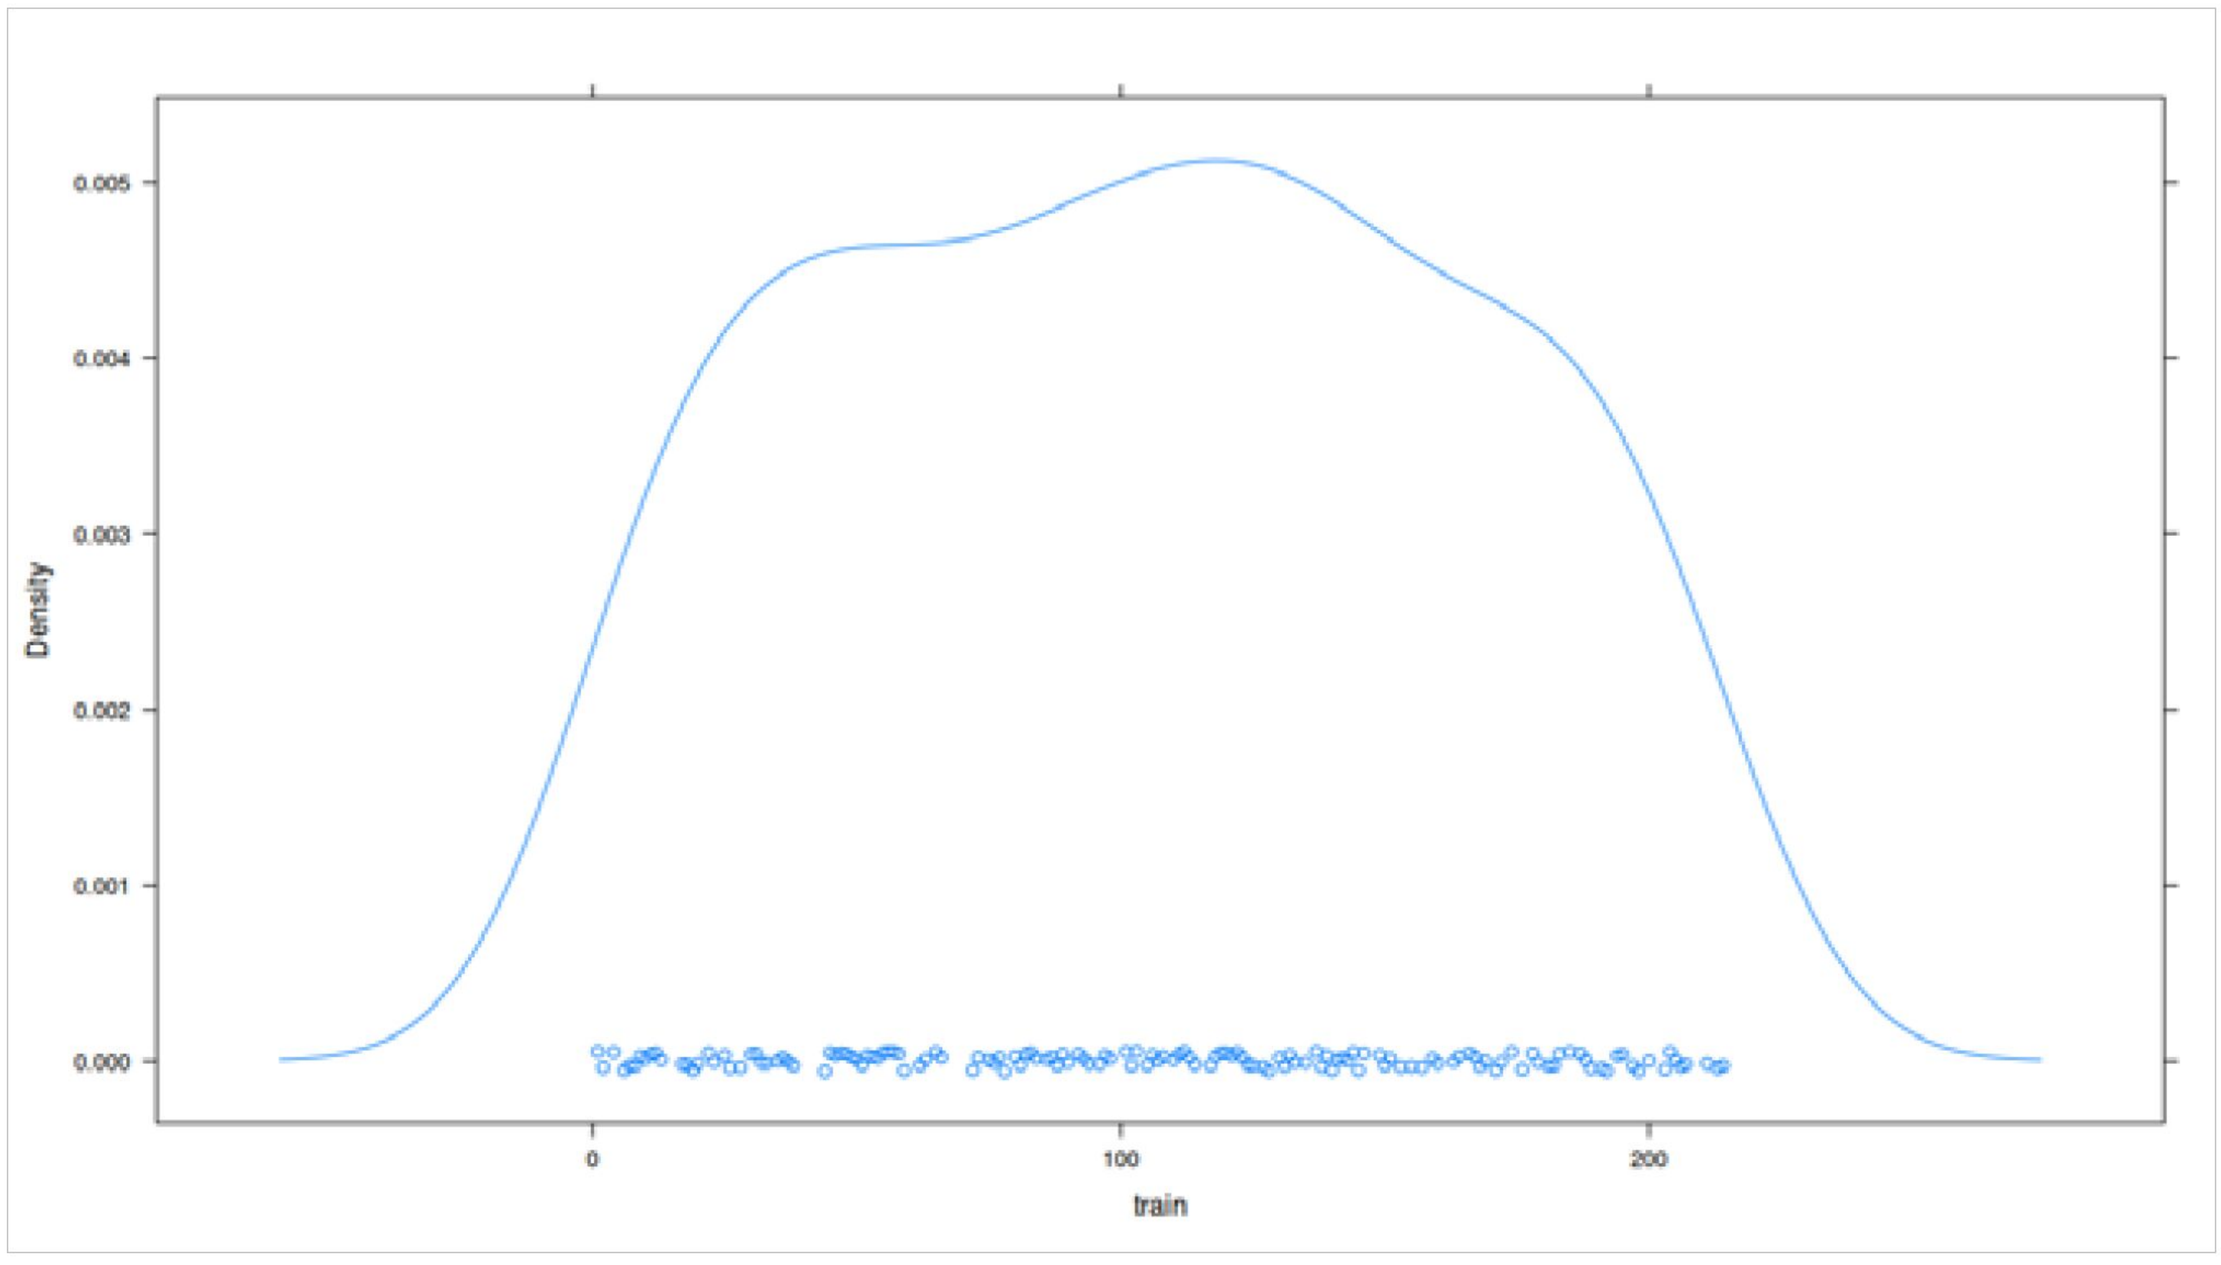


Figure 1: Density plot

```
> set.seed(501)

> fit <- bagFDA(Type ~ ., data)
```

predict()

The *predict()* function can classify observations with *mda*, and the data can be transformed into a matrix using the *confusionMatrix()* function as illustrated below:

```
> confusionMatrix(data = predict(fit, test[, -10]), reference
= test[, 10])
Confusion Matrix and Statistics

          Reference
Prediction  1  2  3  5  6  7
         1 20  6  6  0  0  0
         2  2  9  2  2  0  0
         3  0  0  0  0  0  0
         5  0  1  0  2  0  1
         6  0  0  0  0  3  0
         7  0  0  0  0  0 10

Overall Statistics

               Accuracy : 0.6875
                 95% CI : (0.5594, 0.7976)
    No Information Rate : 0.3438
    P-Value [Acc > NIR] : 2.194e-08

                  Kappa : 0.5757

 Mcnemar's Test P-Value : NA

Statistics by Class:

                     Class: 1 Class: 2 Class: 3 Class: 5
Class: 6 Class: 7
Sensitivity  0.9091  0.5625  0.000  0.50000  1.00000   0.9091
Specificity   0.7143  0.8750  1.000  0.96667  1.00000   1.0000
Pos Pred Value  0.6250  0.6000  NaN  0.50000 1.00000   1.0000
Neg Pred Value  0.9375  0.8571  0.875 0.96667 1.00000  0.9815
Prevalence  0.3438   0.2500   0.125  0.06250  0.04688  0.1719
Detection Rate  0.3125  0.1406  0.000 0.03125 0.04688  0.1562
Detection Prevalence 0.5000  0.2344  0.000  0.06250  0.04688
0.1562
Balanced Accuracy  0.8117 0.7188  0.500  0.73333  1.00000
0.9545
```

skim()

The *skimr* package provides summary statistics for a data set based on the principles of least surprise. You can install and load the library package as follows:

```
> install.packages("skimr")
Installing package into '/home/shakthi/R/x86_64-pc-linux-gnu-
library/4.1'
...

 * testing if installed package can be loaded from temporary
location
 * testing if installed package can be loaded from final
location
 * testing if installed package keeps a record of temporary
installation path
 DONE (skimr)

> library(skimr)
```

The *skim()* function gives a broad overview of the data frame. An example for the training data is given below:

```
> skimData <- skim(train)

> skimData
── Data Summary ────────────────────────
                           Values
Name                       train
Number of rows             150
Number of columns          1
_________________________
Column type frequency:
  numeric                  1
_________________________
Group variables            None

── Variable type: numeric ──────────────────────────
  skim_variable n_missing complete_rate mean   sd p0  p25
p50  p75 p100 hist
1 data                  0             1 105. 60.7  1 53.2
106. 154.  214 ▫▫▫▫▫
```

downSample()

You can use the *downSample()* function to randomly obtain data from a set such that all the classes have the same frequency as the minor class. An example using the oil data set is illustrated below:

```
> data(oil)

> table(oilType)
oilType
 A  B  C  D  E  F  G
37 26  3  7 11 10  2

> downSample(fattyAcids, oilType)
 Palmitic Stearic Oleic Linoleic Linolenic Eicosanoic
Eicosenoic Class
1     9.8   5.3  31.7   51.3    0.8    0.4    0.2     A
2     11.0  5.3  35.0   45.2    1.3    1.3    0.7     A
3     5.9   4.5  24.1   61.7    0.9    0.6    0.6     B
4     6.2   4.0  28.3   59.7    0.9    0.1    0.1     B
...
```

The *downSample()* function accepts the following arguments:

Argument	Description
directory	A character vector of file path names
x	A matrix or data frame of variables
y	A factor variable
list	Boolean value to return a list or not
yname	A label for the class column

upSample()

The *upSample()* function is similar to the *downSample()* function for the fatty acids data where additional samples are added to the minor class as shown below:

```
> upSample(fattyAcids, oilType)
    Palmitic Stearic Oleic Linoleic Linolenic Eicosanoic
Eicosenoic Class
1    9.7   5.2  31.0     52.7     0.4     0.4      0.1     A
2    11.1  5.0  32.9     49.8     0.3     0.4      0.1     A
3    11.5  5.2  35.0     47.2     0.2     0.4      0.1     A
4    10.0  4.8  30.4     53.5     0.3     0.4      0.1     A
...
```

You are encouraged to read the *caret*, *mda* and *mlbench* package reference manuals to learn more functions, arguments and their usage. END

References

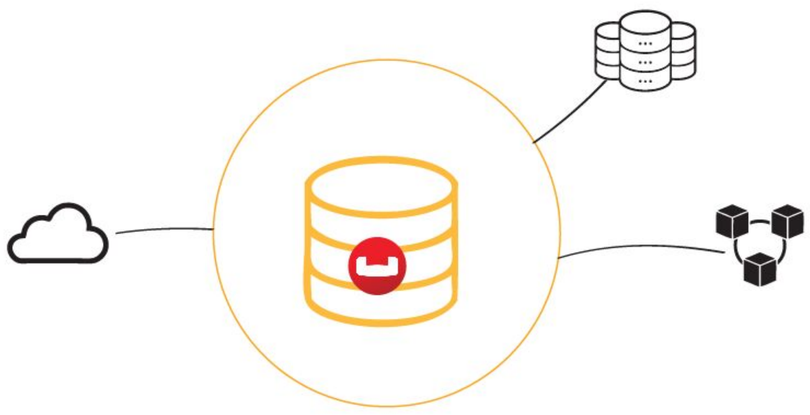
- caret reference; *https://cran.r-project.org/web/packages/caret/caret.pdf*
- mda reference; *https://cran.r-project.org/web/packages/mda/mda.pdf*
- mlbench package; *https://cran.r-project.org/web/packages/mlbench/mlbench.pdf*

By: Shakthi Kannan

The author is a free software enthusiast.

Couchbase: The Distributed Open Source Database

Couchbase boasts of a range of modern features, including memory-first architecture and a powerful query engine. Here's a short introduction to what it offers, with a quick guide on how to install it using Docker.



Couchbase is an award-winning, distributed NoSQL cloud database that delivers excellent versatility, performance, scalability, and financial value for all your cloud, mobile, on-premises, hybrid, distributed cloud, and computing applications.

It is a modern distributed document database that incorporates a built-in managed cache to enable memory-first architecture with all the desired capabilities of a relational database, and more. It exposes a scale-out, key-value store with managed cache for sub-millisecond data operations, purpose-built indexes for efficient queries, and a powerful query engine for executing SQL-like queries.

The services provided by Couchbase are as follows.

- *Data*: Supports the storage, setting, and retrieval of data items specified by key.
- *Query*: Parses queries specified in the N1QL query-language, executes the queries, and returns results. The Query service interacts with both the Data and Index services.
- *Index*: Creates indexes for use by the Query service. The Couchbase Index service supports the creation of primary and secondary indexes on items stored within the Couchbase server.
- *Search*: Creates indexes specially purposed for Full Text Search. This supports language-aware searching, allowing users to search for, say, the word *beauties*, and additionally obtain results for *beauty* and *beautiful*.
- *Analytics*: Supports join, set, aggregation, and grouping operations, which are expected to be large, long running, and highly consumptive of memory and CPU resources.
- *Eventing*: Supports near real-time handling of changes to data. Code can be executed both in response to document-mutations, and as scheduled by timers.
- *Backup*: Supports both the scheduling and the immediate execution of full and incremental data backups, either for specific individual buckets, or for all buckets on the cluster. It also allows the scheduling and immediate execution of merges of previously made backups.

Quick installation of Couchbase server using Docker

Let us look at how we can install Couchbase server using Docker and play around with the Couchbase cluster.

To begin, run the following commands on a console window on your computer:

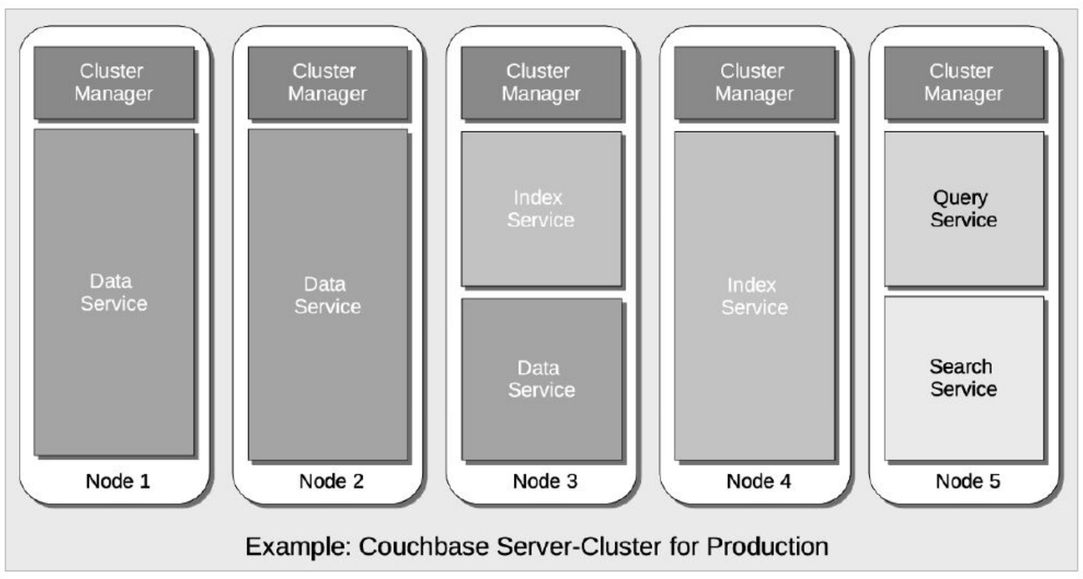


Figure 1: Couchbase server overview

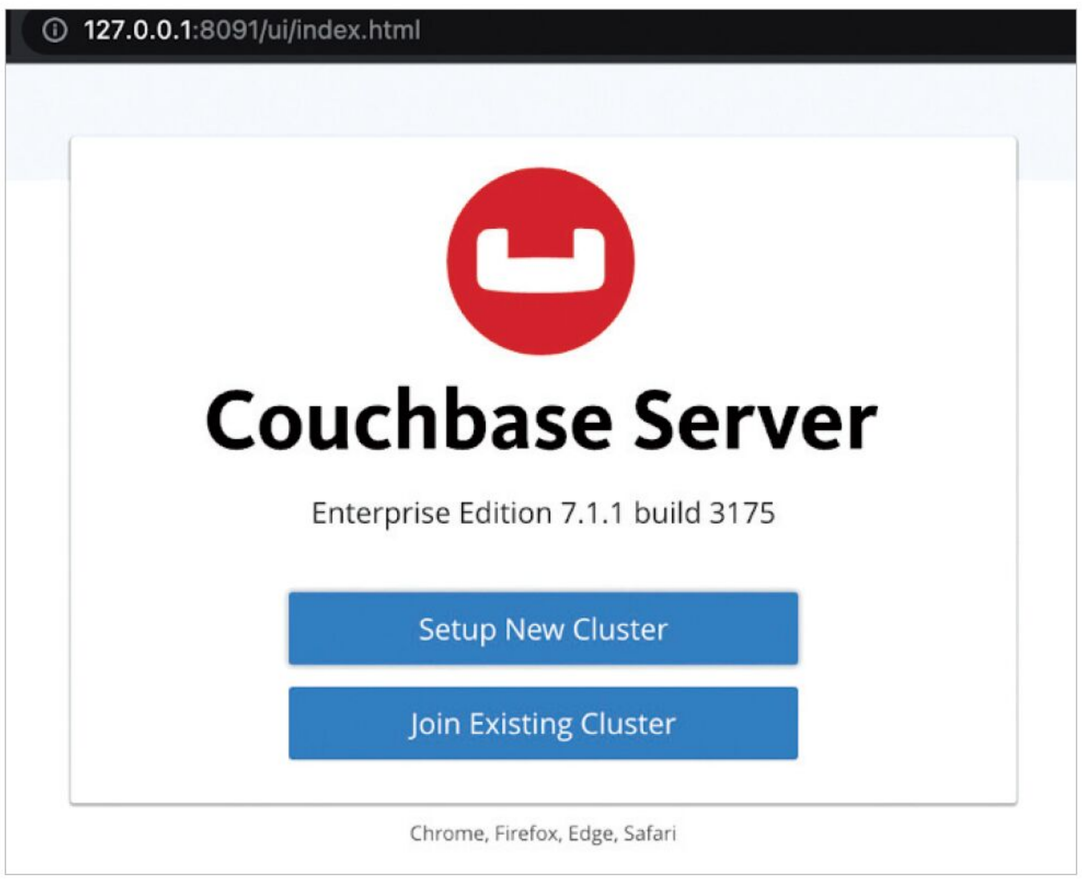


Figure 2: Initial cluster setup page

```
suprithars@Suprithas-MacBook-Pro ~ % docker run -t -d --name
db1 -p 8091-8096:8091-8096 -p 11210-11211:11210-11211
couchbase
docker run -d --name db2 -p 9091-9096:8091-8096 couchbase
docker run -d --name db3 -p 7091-7096:8091-8096 couchbase
```

These commands will help you create a Couchbase cluster with three nodes.

To simplify the process, when getting started, we will provision a single node cluster from the Web console. The Web console can be accessed over the network at *http://<machine-ip-address>:8091/* or *http://<machine-hostname>:8091/*. It can be accessed from the machine on which Couchbase server was installed at *http://localhost:8091*.

You can access the web console via *http://127.0.0.1:8091/ui/index/html*.

Once the above steps are completed, you can click on 'Accept the terms & conditions' and proceed.

In the step shown in Figure 4, you can configure disk and memory settings as well as the services required for your use case. Once the configuration is completed, your cluster will be up and running with one server.

Similarly, you can add more servers/nodes to the cluster with host name/IP address initialised in the beginning (Figure 5).

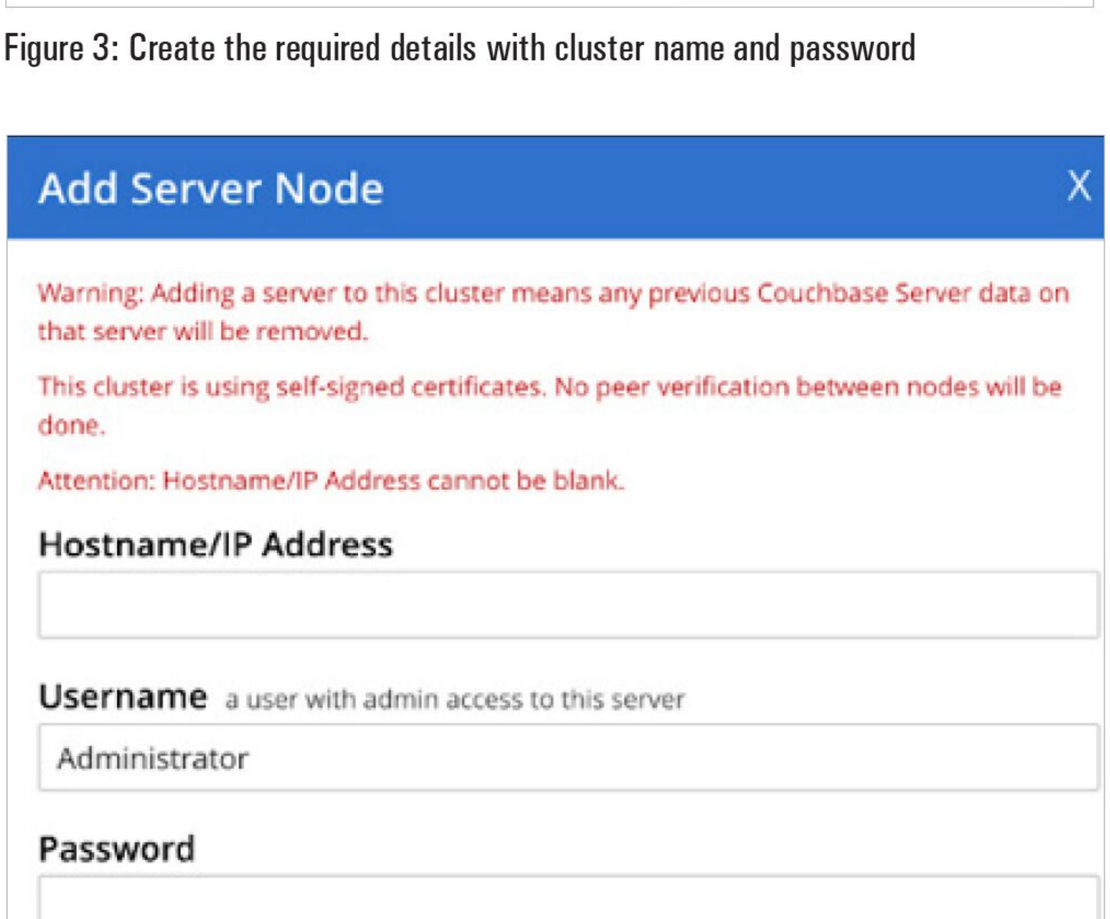


Figure 3: Create the required details with cluster name and password

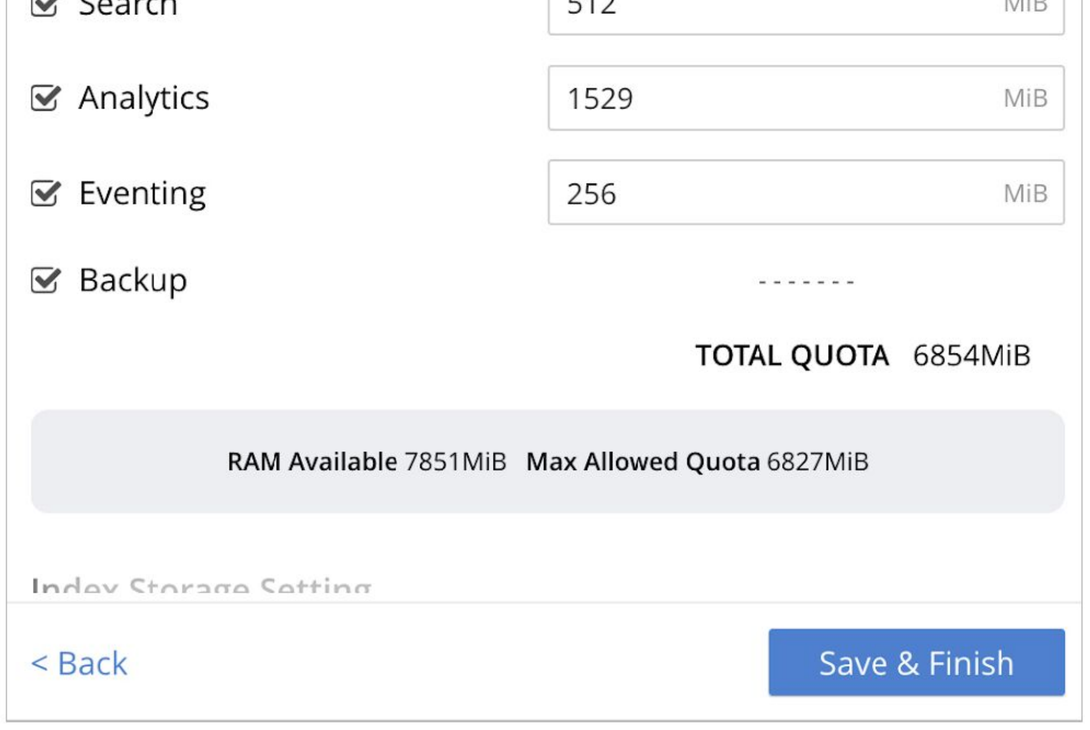


Figure 4: Cluster configuration

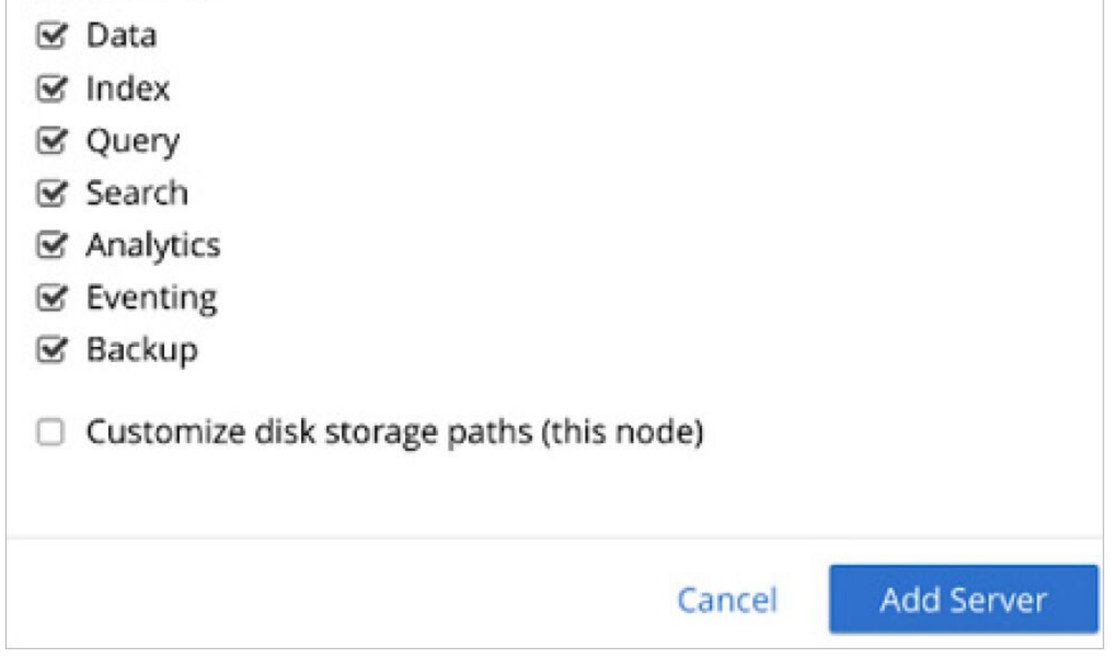


Figure 5: Adding more servers/nodes to the cluster

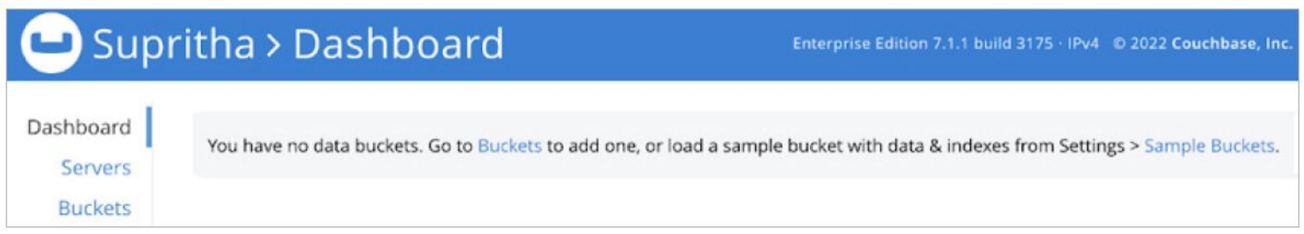


Figure 6: Couchbase dashboard

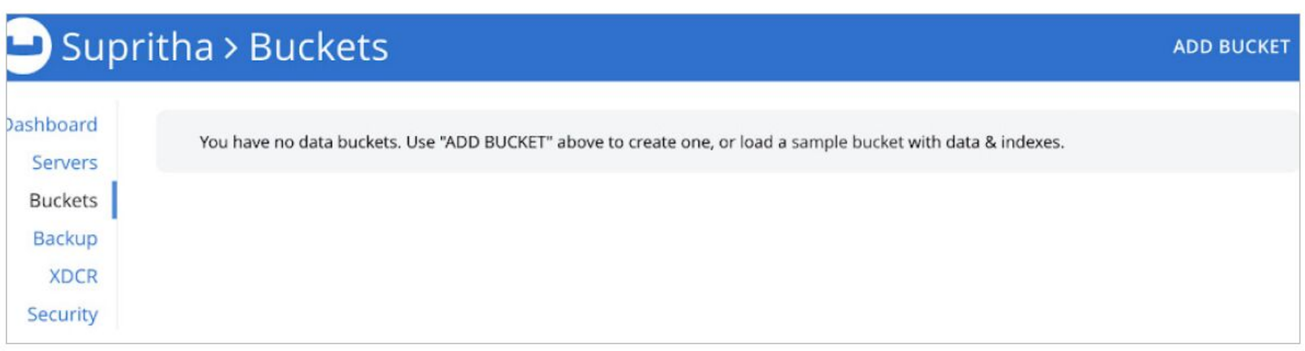


Figure 7: Navigate to *Bucket* tab in the left pane

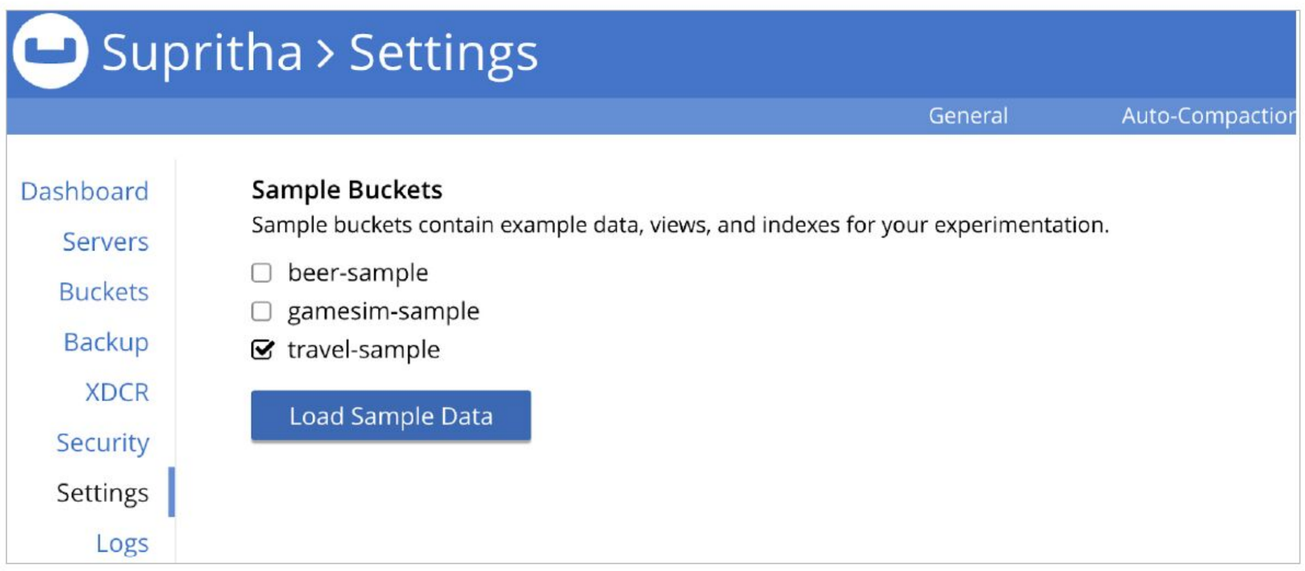


Figure 8: Import sample bucket — travel-sample

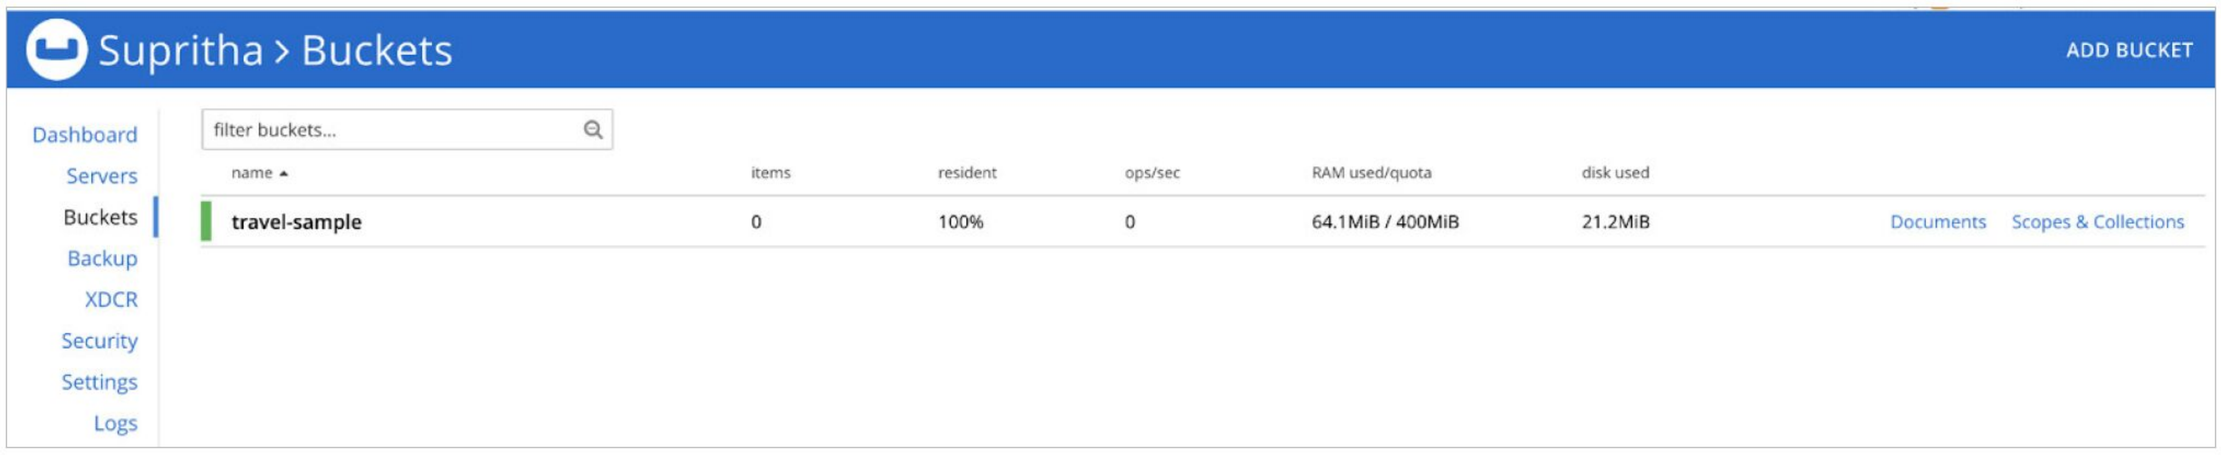


Figure 9: Travel-sample bucket loaded successfully

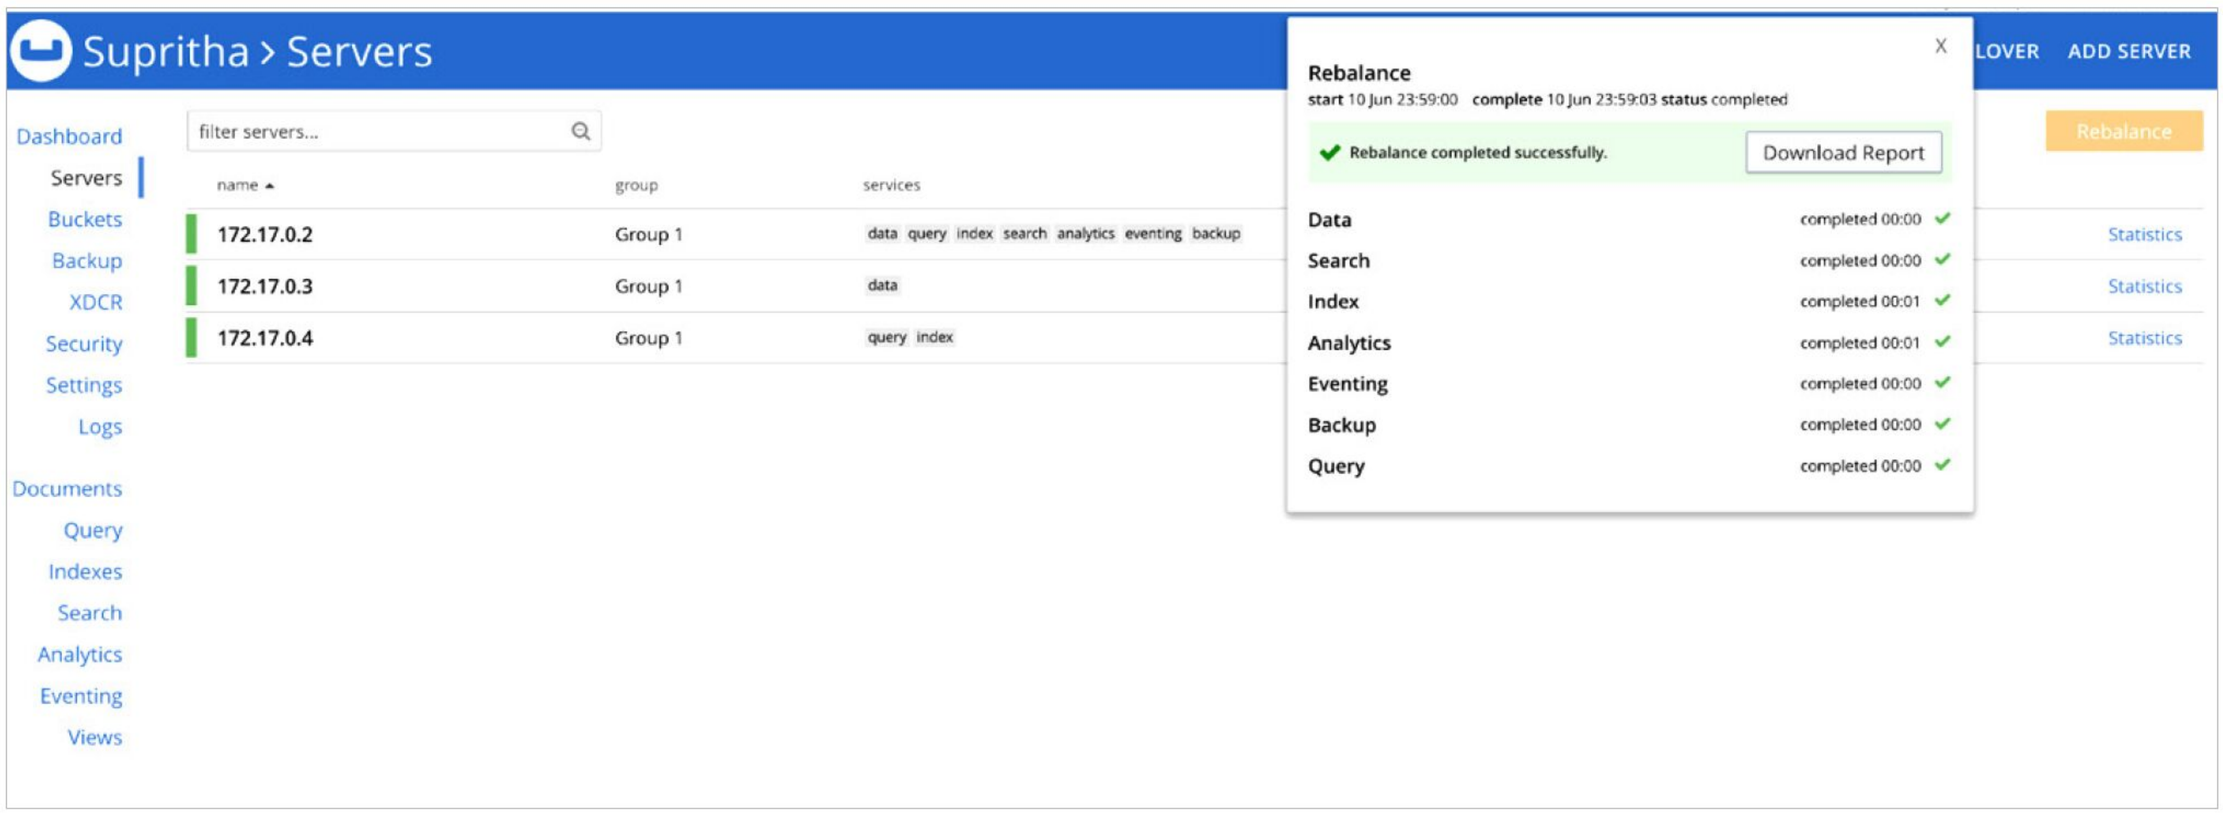


Figure 10: Rebalance

Once all the nodes are a part of the cluster, you can initialise the rebalance. Rebalance redistributes data, indexes, event processing, and query processing among available nodes. When one or more nodes have been brought into a cluster (either by adding or joining) or have been taken out of a cluster (either through removal or failover), rebalance ensures even data distribution. It redistributes data, indexes, event processing, and query processing among available nodes. The cluster map is correspondingly updated and distributed to clients. The process occurs while the cluster continues to service requests for data.

This was an overview of Couchbase with a quick installation using Docker. Make sure to refer to the official documentation for more details and further exploration. END

By: Supritha R.S.

The author is a software engineer and tech enthusiast, driven by a fascination for cloud technology and distributed database systems.

IOTSHOW.IN

31ST JAN - 2ND FEB 2024 • KTPO • BENGALURU

PROFIT FROM IOT



INDIA'S ONLY EVENT FOR THE **CREATORS** AND THE **SOLUTION PROVIDERS** OF IOT

IoTshow.in is India's biggest conference-cum-expo on Internet of Things.

Visitor Registration: www.iotshow.in. *Special Offers for first 1000 Registrants!*

CO-LOCATED WITH:

INDIA ELECTRONICS WEEK

BROUGHT TO YOU BY:



MORE INFO:

Web: www.iotshow.in
Email: support@efy.in
Bulk Registration: iew@efy.in
Tel: +91-9811155335

R N I No. DELENG/2012/49440, Mailed on 27/28th of Advance month
Published on 27th October 2023
Delhi Postal Regd. No. DL(S)-01/3443/2022-24
Licenced to Post without Pre-Payment Licence No U(S)-56/2022-24
Posted at LPC Delhi RMS Delhi 110006

India's No.1 Electronics Magazine
RANKING AS PER ELCINA-IMRB REPORT

"I started reading it when I was a student...
...and I am still reading it, as a student"

—CEO, Design House



electronics FOR YOU
YOURS SINCE 1969

To Subscribe:
https://subscribe.efy.in

OR

Scan This Code

For any query, call: +91-98111-55335 or email: support@efy.in