White paper
Expert Community

What will come next in automation

The convergence of AI, automation and knowledge management will change the way IT Ops work

Introduction

It is often said that IT was sound asleep when automation came to other industries in the past, like the car manufacturing industry. But this is not true! IT not only enabled the mentioned automation in industry and administration but also started nearly from day one to automate more and more of their internal administrative activities in the background, that in the past were done by operators following schedules and work instructions.

More and more of those repetitive, standardized activities where for each event, one predefined reaction was needed to fix one known cause (1:1:1) were taken over by scheduler-software or even hardware, e.g. to replace backup-tapes and store them in tape-libraries.

The truth is that for many years IT automation stagnated, as the same kind of software-based decision-making, which the end-customers would see inside their applications, was not possible in the administration of IT systems as the complexity was too high.

But with recent advancements in technology, it is now possible to handle to a given degree also those events which could have different causes and might need, depending on varying conditions, different solutions (1:m:n). Also, mobile devices, shrinking in size but increasing in capabilities lead a new wave of day-to-day automation. Automation – either visible or invisible - is forming our daily life in the office and at home already in ways thought before to be desirable but impossible.

So, it is no surprise that the questions
'What comes in the next years?' and
'What do I have to do to prepare?'
are today high on the agendas in most companies.

Contents

What is IT Automation

In the five generations model, most companies are currently moving from 2nd generation toward 3rd or some already toward 4th generation tools. Everyone strives to automate (digitalize) as much as is possible and practical. There are some key challenges with automation programs generally.

What to automate

The desire is to automate as much as is possible and practical, the ultimate, a desire being to create fully automated services. However, history tells us this is not always possible or practical. The evolution from low-cost offshore services to lower cost fully automated ones is a key focus area, specifically with the rapid developments in cloud services. It should be an aim to automate at least 50% of non-cloud type services as a starting point given the availability of technology and solutions to achieve this today. The automation of server, database, and application related activities is clearly in focus.

Selection of the correct platform & tools

Finding a platform and technology solution that will deliver the promised and anticipated benefits over a future term is a key challenge, solutions come and go and change fairly rapidly. A "one size fits all" is probably not appropriate for large enterprises. Flexibility and attractive license terms are almost as important as functionality. A very careful assessment of the needs of the organization must be considered in the procurement cycle. A multi-vendor approach may be most appropriate when constructing an automation platform. History has taught us claims made by vendors are not always fulfilled. A layered platform/tooling approach may be most appropriate. The introduction of new and emerging technologies into the automation arena, e.g., Cognitive or AI also adds to the dilemma.

Availability of staff to build, test, implement and maintain the automation

After having selected the appropriate platforms and tools, the key and major challenge is the organizational ones. The organizational approach to the investment of technical staff to deliver the end to end automation must not be underestimated. Larger organizations often fail to deliver the expected benefits of automation due to lack of organizational alignment, ownership boundaries and investment of key people. Strong management governance, leadership, and planning are required to achieve good results, failure to consider this will compromise future results.

There is also a significant organizational consideration regarding the placement of automation skills within the company; there are two schools of thought typically in this area: -

- Have a separate and dedicated automation team
- Have automation teams and skills included as a core function of each services team, e.g., SAP, database, server, workplace, etc.

There are pros and cons to each approach, future automation requirements, and the drive to automate as much as possible to reduce costs and increase competitiveness being at the heart of the selected approach. Given that in the next 3 – 5 years the outlook will be for the most modern infrastructure services being fully automated a careful review and implementation of the appropriate structure is paramount.

Availability of platforms and the infrastructure

It may seem a fairly obvious requirement, but having access to the customer's applications, services, and infrastructure components to enable automation to be implemented will often lead to delays in delivery of an automated solution. The requirement to be able to build fully, test and then implement automation requires possible outages to services, or waiting for the next available scheduled maintenance window to do this, issues arise, and delays will inevitably occur during implementation. Careful planning is required to mitigate such issues. The adoption of DevOps can be a helpful enabler in this scenario.

What comes next

Combined with Artificial Intelligence (AI) the automation will soon make the next big step, toward Machine Knowledge and even Machine Responsibility. What brings a new challenge when IT services rely on cognitive automation: Can you trust a machine decision?

For such a system to work it needs three main processes:

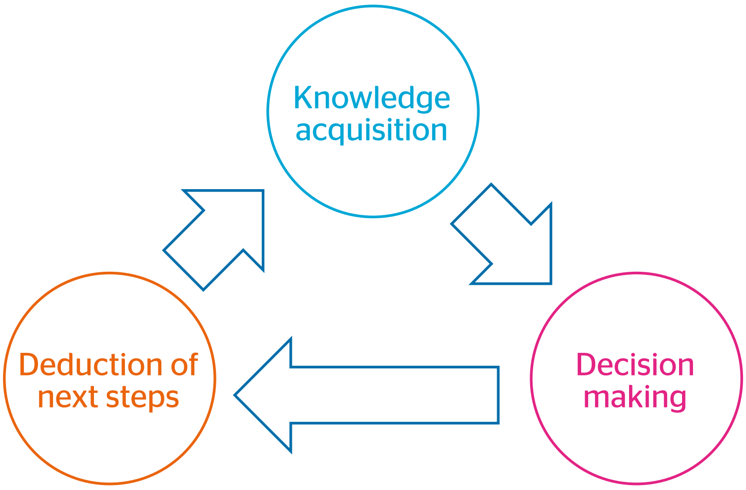

Knowledge acquisition:
The ability to collect and aggregate relevant knowledge from various sources.

Decision making
Use of the available knowledge to

- a) recognize that a decision needs to be made,
- b) to formulate a decision-making matrix,
- c) to reach a decision and
- d) to evaluate the correctness of the decision

Deduction of the necessary next steps
The ability to produce a workflow of activities that are tailored to the decision and current situation and to start executing it.

Beyond the AI system used for the decision-making and deduction of necessary steps, the knowledge base (KB) is the other fully-interlinked key component here. It has to store all the complex structured and unstructured hard information and rules to be used by the AI component in the decision-making process. The process itself is also predefined by the content of the KB and the allowed rules.

AI can work only based on the quality and quantity of data available in the KB. While this content can be checked, and AI tested intensively, there will be sooner or later the point where the KB-induced limits are reached. This limit can be only overcome by expanding the KB to give the AI additional possibilities to find (new) solutions to (new) problems. In a fully automated solution, this expansion of knowledge would be automated as well. The IT industry follows here two different approaches:

IBM's Watson is designed to collect its knowledge self-reliant and independent from a large number of sources, curated and un-curated, in the 'quantity over quality' approach also used in Big Data. While this offers unprecedented access to the knowledge at all levels, it also comes with the risk, that the KB might be contaminated with wrong facts that could influence the decision-making process and lead to wrong decision making. Simpler systems using this open data collection process, like the 'Tay'-Twitter Chatbot from Microsoft has already fallen victim to such contamination and turned to become a 'Hitler-loving sex robot'[1]. So, sources of knowledge may willingly or unwillingly have been tampered with, which may lead to disaster instead of cure.

After this, the IT industry turned more to a second approach, the usage of curated knowledge in a machine-readable form. The knowledge base could turn to some trusted sources and learn from them what it needs to solve a current problem. Software vendors could make their knowledge about their products and errors available to help the automation system for their customers in solving issues. This 'quality over quantity' approach is safer with regards to contaminated information, but of course, is barring out all additional knowledge that is not yet part of the knowledge published by the trusted sources. And of course, such trusted sources also could be hacked, and the content so is contaminated.

The 'quality' of knowledge not only relates to the fact that it doesn't contain malicious content, but also that the problem for which this knowledge should be used is precisely described.

So you can use the knowledge that is more specific for the issue at hand. This helps to avoid more general solutions, that might fix your issue as well, but have a way larger impact. One example for this is that a large number of people always recommend booting a windows system if an error occurred because they don't know the cause of the offending error, and so hope that the cause might be resolved when the system is up and running in the restarted-configuration.

A reboot indeed might help to solve the issue but restarting only the offending failed service might have solved this issue as well, without the downtime and all related implications for the usage of this system following a restart. The more specific the knowledge is, the faster and impact-reduced the solution can be.

http://www.telegraph.co.uk/technology/2016/03/24/microsofts-teen-girl-ai-turns-into-a-hitler-loving-sex-robot-wit/ [1]

What comes next

In the IT business world, the best sources for good quality knowledge are the vendors of the products. They know the ins and outs of their products. So, for the automation systems to use such vendor KBs, they must come in a machine-readable, standardized format. The way such an "XML for knowledge" and the interfaces to such sources would look like, how they are billed, etc. is still under development and ongoing discussion. The interface topic might be less of an issue when knowledge comes via an organized process by a vendor but is more challenging when knowledge from the web with all its kinds of formats (text, video-tutorials, flash-animations) and sources (Websites, discussion-boards, twitter-feeds) come into play[2]. Also, the way to organize processual knowledge in a way that it can be described and how the automation systems raise their questions to the knowledge-base is still not clear. As the industry always needs to have standards, general-use technologies like GraphPath[3] are developed.

With the ability to expand own knowledge autonomously from external sources, a big step is achieved to hand over responsibility to a machine. This autonomous expansion can be on a regular, planned basis (e.g., looking for updates every day in all sources) or on demand. Most likely we will see a mix of both.

New knowledge is not only needed to fill gaps in existing knowledge, triggered by the knowledge-gap procedure, but also to improve existing knowledge, triggered by the update procedure.

Additional features are needed to reach the same responsibility level as a human: The ability to review its work and to call for help when stuck, or unsure which way to turn (technical escalation). The latter is best practice today but might also enhance the ability to review own work and spot potential errors in its decision-making process. Of course, each decision -making AI is designed to make as few errors as possible but rechecking every step and decision along with a decision tree by default would overload most AI systems quickly and take away any speed advantage, as, in the end, AI would check everything to verify every possible condition. Such behavior in most cases is neither warranted, as the AI would be correct to an increasing degree in what it did nor helpful for the system as such, as it would make the AI too complex. The preferred solution would be a split of work. One AI system does the normal work, while another can be contacted to review and evaluate the work done. The review functionality and based on that the process to suggest improvements and fix errors could also be handed over to a machine before a human needs contacting when this second system also comes to its limit.

All this is in place already to a limited degree, but with progress on an already exponentially increasing learning curve, but still a business secret in most cases.

The last area associated with the responsibility is accountability[4]. Our human ethics and judicial systems always call for somebody to be accountable for intentional actions. This decision is already in full swing on all levels of the problem. Two well-known examples of this are 1.) Alexa systems ordering products for their owners simply because they heard somebody do this on the TV or radio[5] and 2.) decision errors of autonomous cars[6], which misidentified input or misinterpreted internal rules. These examples touch the topic of intended/unintended action, like when a human moves the arms during a nightmare while asleep and trashes something on a nearby night table. (Here a 'known unknown' came into play; the interaction with other systems, beyond the control of the designer, but now is a 'known,' see below in this article.)

This legal and ethical topic will stay with us for the foreseeable future as there is always room for such errors. As long as the rules along which decisions are made, and the knowledge on which they are based is in the full responsibility of a manufacturer, he might be accountable[7]. But with self-learning and self-developing AI and their autonomous expansion of knowledge, especially when coming from un-curated sources, the contributions of others might become bigger and more important. In the end, it will be as difficult to judge who is responsible as we also see in human crime cases with a large number of persons that contributed to a situation on a causal level, but didn't collude. But this will not stop the technological process, humankind is used to develop its legal system in parallel to technological and social progress and accepts that at first things might not be punishable until laws or better laws are in place.

At the other end of the spectrum of this development, there are already discussions who would/could get an award for intellectual property or scientific milestones when AI systems made this progress possible, e.g., interlinking extended fields of knowledge like Watson is doing.

Ultimately all this might lead to machine ethics, which goes way beyond the famous three laws of robotic of Isaac Asimov[8].

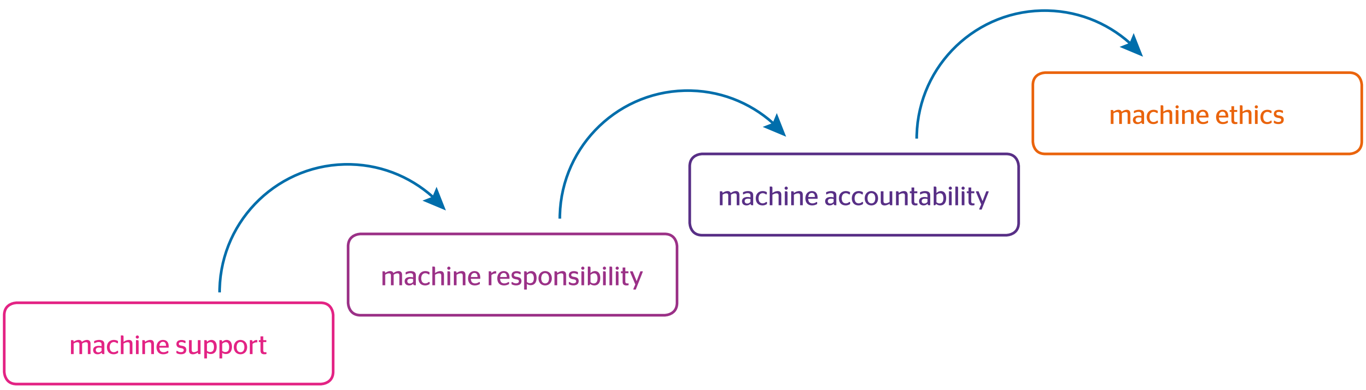

[2] https://www.linkedin.com/pulse/machine-readable-ontologies-semantic-web-randall-shane/
[3] https://techcrunch.com/2017/08/23/graphpaths-knowledge-graph-as-a-service-could-insert-ai-into-global-corporates/
[4] https://www.diffen.com/difference/Accountability_vs_Responsibility#
[5] https://www.theverge.com/2017/1/7/14200210/amazon-alexa-tech-news-anchor-order-dollhouse
[6] https://www.scientificamerican.com/article/who-s-responsible-when-a-self-driving-car-crashes/
[7] https://cyber.harvard.edu/publications/2017/11/AIExplanation
[8] https://en.wikipedia.org/wiki/Three_Laws_of_Robotics

How to prepare for this

All this - of course – comes with some necessary changes in the way we work and learn to work in the future.

Technology assessments needed

Predicting the (technological) future is always difficult and error-prone[9], but ignoring issues in the hope that the benefits are outweighing the disadvantages is not a good way forward. Good management calls for technology assessment, which forms the base for further investigations and developments to mitigate identified risks.

Example 1: The automation will raise the level of process-driven work to new levels. What is good from an economic point of view also raises the problem that there are always cases that don't fit into the design of the automation, but still need to be handled. Such irregular cases will lead to errors and complaints about unfair decisions. Sol solves these issues special processes, and maybe additional human teams are needed to handle them. These teams must have the knowledge how such exceptional cases still can be processed correctly and how to review and correct unfair, automated decisions. These efforts can cause a conflict with the hope that automation will reduce the need for human intervention and that it can reduce the amount one needs to know about a system to operate it. As nobody can know everything, and existing documentation is often not sufficient, islands of expertise for one or another part of a process or the related tools will develop again. The need to have people with special knowledge in sufficient numbers and full documentation to handle exceptions are often underestimated or ignored as this will decrease the potential business case for automation.

Example 2: The dependency on those systems will become very quickly very big if the workforce is scaled back the way the pre-calculated 'automation savings' request it. With very high automation levels the size of the remaining workforce might not be big enough to handle even the incoming tickets (very) high priority (P1/P2) during the downtime of automation. While at the same time the customers will escalate more and more tickets that have a normal or low priority (P3/P4), as they very quickly will get used to very short resolution times also for P3/P4s and are less willing to wait.

While the simpler automation of the 1-3rd generation was still easy to manage when it comes to risks, the combination of two complex and increasingly powerful technologies (AI and automation), will require new technology assessments for each domain that is planned to be automated. For this typically the four quadrants of the Technology Assessment is used as a starting point.

What we know that we know	What we know that we don't know
This knowledge is reflected in the specifications, handbooks or the software itself, plus some not documented knowledge in the heads of the designers and engineers.	Typical examples are software or algorithms that are used, which are the not-open source but business secrets of other companies, but also known assumptions instead of proven facts or 'black box' mechanisms used.
What we don't know that we know	**What we don't know that we don't know**
These are typically unrecognized rules which were included, e.g., unintended bias in algorithms like Google search or recognition-algorithms that ignore the possible diversity of bodies or objects they have to handle[10].	Everything that is not yet understood to be a factor or risk in this context or facts and rules assumed to be correct that is later to be found incorrect.

An important point that has to be reflected here is the ethical limitations that should be observed by automation. On example here is the ethical dilemma in autonomously driving cars that might have to decide in an accident situation whether it is acceptable to hit pedestrians to protect the lives and health of its occupants[11] or not; other examples are arranging price-riggings[12] or cheating in the stock market[13].

Based on the risk analysis there might be some research for the specific use case necessary to reduce the unknown facts, even when one never will be able to reduce these unknown risks to zero.

http://rodneybrooks.com/the-seven-deadly-sins-of-predicting-the-future-of-ai/ [9]
https://mic.com/articles/124899/the-reason-this-racist-soap-dispenser-doesn-t-work-on-black-skin#.OkGqKzcT1 [10]
https://www.theguardian.com/science/2016/jun/23/will-your-driverless-car-be-willing-to-kill-you-to-save-the-lives-of-others [11]
https://www.newyorker.com/business/currency/when-bots-collude [12]
'Dark Pools' by Scott Patterson and 'Flash Boys' by Michael Lewis [13]

How to prepare for this

Human-Machine Collaboration

While Machines will be increasingly able to take over some works completely, these activities are normally just a part of what humans do. Machines still have a large problem in making decisions in complex settings or where they don't have predefined rules along which to make the decisions. Most studies agree that Machines are not going to replace humans completely but will change the job profiles significantly, and so we will see collaboration between humans and machines on a new level.

To make this collaboration work, the employees have to learn how to use the tools and applications that are coming with the mix of Automation, AI, and Analytics (A³).

The focus of the human activities will here be on ensuring that the data the A³ uses are available, correct and complete, supervise the automation and review that the results are acceptable, potentially results are to be reviewed before execution and make decisions if the machine can't do it. This approach is recommended even for already more mature appearing systems as components used by A³ might fail or be disabled for one reason or another and so make other than the expected decisions or no decisions at all[14]. Until confidence into and understanding of the to be expected outcome is big enough on the user side, it might be necessary to integrate repeated check-points where the automation is presenting its current results and planned next steps to its user for review and confirmation or abort if necessary. While this might reduce the efficiency of automation at first, it will be necessary during the phase where the users are learning how to use the automation and protect everyone from possible user-errors or automation that is producing unwanted results and side-effects. This work-sharing allows both humans and machines to be used in the best possible way.

Training

To make the new automation-based tools accessible to end-users, new applications and interfaces have to be built and maintained. At this point, automation becomes the main task of an IT operator. We finally reached DevOps. A highly automated system needs highly sophisticated maintenance. Therefore, operators must be trained early to accomplish a future task while the automation is doing the hard work. But the main challenges are the building of A³-systems.

The first activities that are automated are normally simple/easy tasks. While no engineers will miss them in their daily-life and often is happy that they were automated these are often also the first tasks we let beginners do as part of the learning-by-doing approach often used in IT. With these tasks disappearing from daily operations the bar, over which new employees must jump when they want to work as an engineer in the delivery, will rise quickly. Therefore, we must find ways to compensate for these 'lost first steps' in their developments. At the same time, we must keep the existing engineers also fluent in these basic tasks as the automation might suddenly fail so that manual execution and so the knowledge about these 'disappeared activities' is necessary again. So, we will need 'simulators'/ labs to train these skills, like what pilots are doing to stay fluent in all tasks and keep their license.

Automation will bring changes to education as well. Aside from traditional learning students and employees will need to develop different skills such as adaptability to the change and self-learning capabilities.

Experience has shown that the approach building automation by piecing together several contributions from experts with a deep but narrow scope into one bigger workflow is time-consuming and error-prone as a lot of gaps can be found in those workflows and so automation fails. The better approach is where one automation expert takes responsibility end-to-end and thinks through the whole process and gather most of the needed knowledge oneself, consulting with the deep experts. Such an approach will require a deep understanding of the involved automation tools, their APIs, and way of working. It will also lead to a change in the self-image of those engineers; they are no longer OS, database or application experts, but more generalists with a broad but not so deep skill over many technologies again. After a decade where deep specialization was the key to a career in IT, HR will have to open career paths with corresponding training curricula that reflect this new skill-demand. We will have to find those people who are willing to work on such a broader scope of technologies again. Such a new role requires permanent learning on a much broader range of skills than before as a deep-expert.

Human Resources needs to keep pace with the changes and look for candidates that will bring value to the company. The technical skills for the job are still necessary, but as well a new type of profile that can be more focused on abilities such as creation, organization, supervision, leadership, change adaptability, think beyond goals, emotional intelligence, and self-learning capability. All these abilities might not be easy to develop on current employees, so it's important to the employers to implement creativity workshops, self-training programs, cross-functional team meetings and brainstorming on different topics looking always to grow the business.

High-level education might need to change regarding content and length. Traditional degrees take around four years to complete, but in the future, that might be too long to force students to adapt fast to the changes and educate more with focus on leadership, organization, and creation of jobs. At the same time, it might be that the wide-spread automation makes skills needed in the past less important, so it can be skipped to compensate for the new content.

[14] https://www.azcentral.com/story/money/business/tech/2018/03/22/what-went-wrong-uber-volvo-fatal-crash-tempe-technology-failure/446407002/

How to prepare for this

Documentation

Even the architects and designers might have difficulty to understand the current automation, especially when they are combined with AI like in Arago Hiro or similar tools. Evolving them further without causing unintended side-effects will be especially hard if there is no full documentation of how everything is done, how all those things play together, which data and rules are used for automated decisions and interact. The risks of unintended interactions and consequences are raising, not only in volume but also in scale/impact. While it was okay for decades to build and use a "black box" in the machine learning environment, engineers, customers and legislators are calling more and more for "transparent boxes." When using a "black box," it is not possible to retrace step by step how and why the AI decided as there are no records based on which data-points the decision was made and which rules were used and why. In "transparent boxes" these records are available, at least of the request in an "observation mode" to have a full understanding of how the AI is coming to its conclusions and decisions. More transparent systems would also fulfil the request to have the possibility to audit the methods and factors the AI/automation uses to come to its decisions, and make corrections if needed. All this is especially important when automation must interact with each other and not only handle a restricted, small task as interactions between AIs can cause a lot of unintended behaviours and results[15].

Processes

To make automation a success also the processes in the companies have to change, else the speed-gain that is coming with automation might be lost, at least partly. One typical example from an end-user view might be the approval process for an order. In the normal processes, such an approval workflow can take days to be complete. In an automated world, the manual checks and approvals done by managers and controllers can be automated as well. The automation can check whether a user is eligible to order this product, whether funds are budgeted for it, and approve automatically, requesting human interactions only above a given value or when refused.

While the changes in the event- and incident-management are often apparent and well understood, the increase of complexity will require changes to change and problem-management that are often overlooked or under-estimated. Especially in highly interactive, self-learning automation environments adding new features, possibilities and knowledge can cause unexpected side-effects. These risks must be taken seriously in change-management and fully analyzed in problem-management or the whole IT environment might soon become a black-box and is no longer under control.

> Like humans, self-learning systems, especially when supported by AI might use new features in an unexpected or unintended way.

Knowledge-based AI is trained to find the most efficient way to solve the problems. Such freedom in recombining knowledge and data means it will test from time to time whether the currently used path is still the most efficient one and so tries to recombine knowledge to find better ways. Such behavior will happen when new possibilities in the form of new knowledge, features or functions are added. These tests are part of the permanent improvement function of self-learning and are to be expected. It is recommended to give these systems learning and testing environment, where it can learn to use the new knowledge and functions in the best way before it can use them on systems in the production environment. For changes, it might become necessary for manufactory to run such tests first to make a risk assessment for the planned change. Without such assessments, it will be very difficult for change-managers to understand and rate the risks.

These tests and simulations are also vital for the Problem Management process when errors that occurred in the production environment need to be analyzed. Full documentation in log files of every data and rule used by the automation might be impossible due to the scale of systems in production. So test environments are needed, in which due to the limited amount of systems/transactions such detailed logs can be done. These logs then can be analyzed by problem-managers, designers and automation experts to learn how their current and future system installations might behave and why. Only then they can fix errors. Else, it would remain a black box, which is not acceptable, given the importance and dependency on such systems.

http://www.dailymail.co.uk/sciencetech/article-4747914/Facebook-shuts-chatbots-make-language.html [15]

Conclusion

IT automation will enable an increasingly strong collaboration between humans and machines. Machines will soon be able to learn and expand their knowledge and rules-sets autonomously to close gaps in their existing knowledge base and so even learn new activities.

To be able to use and manage the self-learning systems of the 4th and 5th generation of IT automation, companies must adjust the way their IT departments work.

- Define digitalization mission and vision at the executive level and communicate as digitalization roadmap from top to down,
- Adapt the digitalization training curriculum for the employees in general and form in the IT department IT generalists again, who can design and manage automation end-to-end,
- Build labs/simulators to keep the needed knowledge for automated activities alive in the minds of the employees for a time when the automation failed,
- Possibilities to simulate the automation under real-world conditions to get a better understanding of changes and problems and to avoid unintended consequences/interactions,
- Up-to-date, detailed documentation to be able to understand how the self-evolving systems work,
- Adapting the company internal processes to use the benefits of automation fully.

About the authors

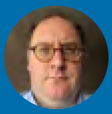

Matthias Noch

Born in 1972 in Hamburg, since 1990 in IT, in automation since 1998, with Atos since 2006, 2010 Global Service Architects, 2014 PoC for Arago Hiro and involved in automation and AI since. Automation Subdomain leader for Intelligent Automation.

Klaus Halbig

Born in 1967, in professional IT since 1991, automation since 1999, with Atos since 2012, current role Global Service Architect Tooling.

Tony Hand

Born in 1958, been in IT since 1978, recent roles 2005 – 2010 Head of Workplace services UK, 2010 – 2015 Global Head of service architecture workplace, 2015 – present IDM SWAT director.

About Atos

Atos is a global leader in digital transformation with over 110,000 employees in 73 countries and annual revenue of over € 11 billion.

European number one in Cloud, Cybersecurity and High-Performance Computing, the Group provides end-to-end Orchestrated Hybrid Cloud, Big Data, Business Applications and Digital Workplace solutions. The group is the Worldwide Information Technology Partner for the Olympic & Paralympic Games and operates under the brands Atos, Atos Syntel, and Unify. Atos is a SE (Societas Europaea), listed on the CAC40 Paris stock index.

The purpose of Atos is to help design the future of the information technology space. Its expertise and services support the development of knowledge, education as well as multicultural and pluralistic approaches to research that contribute to scientific and technological excellence. Across the world, the group enables its customers, employees and collaborators, and members of societies at large to live, work and develop sustainably and confidently in the information technology space.

Find out more about us
atos.net
atos.net/expert-community

Let's start a discussion together

For more information: matthias.noch@atos.net

CT_j2294_190411_RY_WP_WHATWILLCOM