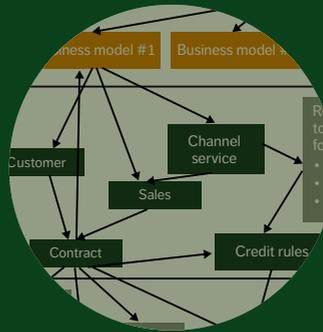


Review

ERICSSON
TECHNOLOGY



AI-NATIVE BUSINESS SUPPORT SYSTEMS



ERICSSON

BSS and artificial intelligence

–TIME TO GO NATIVE

The growing need to support disruptive services emerging from the Internet of Things (IoT) and 5G requires a fundamental transformation of business support systems (BSS). At Ericsson, we believe that the best way to achieve this is by forging BSS and artificial intelligence (AI) together to create truly AI-native BSS.

LARS ANGELIN,
JOHAN SILVANDER

Although AI is of obvious benefit in terms of business optimization, and has been used in all sorts of businesses for decades, AI and BSS have never been integrated into one efficient system.

■ Examples of areas in which AI is already used in conjunction with BSS software include customer retention, chatbots, revenue and cost predictions, customer analysis, customer experience management, customer yield optimization, automation, process reengineering, simulations, quality improvements, and fraud and anomaly detection.

AI capabilities enable improved business decision dynamics and better decision precision, resulting in better business performance and agility.

Virtually all business activities can and will benefit from AI, and as 5G and the IoT continue to expand, the number of use cases will only continue to grow. The challenge we face at present is that the learning, insight-building and reasoning capabilities of AI in today's telco BSS are not as strong as they need to be to cope with emerging use cases.

For the most part, AI capabilities today are simply bolted onto telco BSS one by one. But this is inefficient in terms of life-cycle costs, because the BSS must be repeatedly upgraded to benefit from the AI algorithms. Further, as they are separate systems, the BSS information must be transformed to fit AI systems, and vice versa. A much more efficient alternative is AI-native BSS – that is, BSS with intrinsic AI capabilities where the AI logic

is a natural part of BSS logic in terms of both design and operation. This approach results in a system that can handle more complex business situations, generating more optimized business outcomes.

BSS evolution drivers

The main growth opportunities for communication service providers (CSPs) within the next decade are 5G and the IoT, with an estimated annual value of approximately USD 600 billion [1, 2]. To capitalize on this opportunity, CSPs must be able to support a marketplace with an ecosystem of many actors that have their own business models, where each actor may be both supplier and customer to other actors. Business support complexity increases dramatically in this environment. Current telco BSS, which can only support a single enterprise shop with a few business models, are simply not up to the task.

Surveys show that a clear majority of the telecommunications industry actors expect AI to have significant business impact in the coming five years, affecting both the top and bottom lines. They also expect AI to bring a significant competitive advantage to the enterprises using them, growing proportionally with AI usage. Analysts predict that enterprises will invest in AI competence, AI maturity and in organizational AI capabilities [3, 4], despite the costs [5].

ARTIFICIAL INTELLIGENCE

Artificial intelligence (AI) depends on software algorithms. At Ericsson, we use the term AI in its widest sense, including several subfields such as machine learning, representation learning and deep learning. AI-related areas such as natural-language processing, automated reasoning, multiagent systems, symbolic learning, knowledge representation, intelligent tutoring systems and high-level computer vision are also included [8].

There are essentially three main forces driving the combined AI-BSS evolution [1]. Firstly, business agility is a highly valued BSS property since the business itself evolves and new business opportunities emerge. AI plays a key role in both identifying opportunities and in shaping the new business models to pursue them. Secondly, the maturity of the communications industry is driving ever lower business transaction costs, as demonstrated by existing platform players like Amazon and Alibaba. In light of this, AI-supported process automation and reengineering are the tools

Terms and abbreviations

AI – Artificial Intelligence | BSS – Business Support Systems | CRISP-DM – Cross-Industry Process for Data Mining | CSP – Communication Service Provider | IoT – Internet of Things

●● A BUSINESS INTENT STATES THE DESIRED OR OPTIMAL OUTCOME OF A GIVEN SITUATION ●●

of choice. Finally, the cloud provides an ideal foundation for continuous introduction of AI and BSS marketplace capabilities. This is because the cloud offers deployment flexibility, elastic scaling and a micro-service architecture that enables a more fine-grained separation of concerns and componentization with loose coupling.

Key challenges to adding AI to BSS

Introducing AI into existing BSS is not straightforward. Some challenges, such as data acquisition, data piping and training algorithms, are obvious and well known. Others, however, are less so. One example of a less obvious challenge is the fact that the interpretation of the AI results requires business competence; another is that monetization requires both retraining within the organization and redesigning of the existing set of business rules and processes and system reengineering [5].

Traditional, non-AI-native BSS are divided into component silos – such as customer relationship management, catalogs, billing and order management – each with their own information. This arrangement contradicts AI efficiency and dynamics enablement, which require an open, pan-BSS information and rules view. When an AI capability is added to this environment, it is treated as an add-on, requiring both AI and BSS system competence, information transformation and in many cases partial system re-implementation or reconfiguration. BSS performance issues such as latency and scaling may arise.

Many business situations are multifaceted, have many root causes and may include both gains and risks. In many cases, a main business intent (also known as a KPI) must be broken down into a combination of subintents. These will be based on many data sets and algorithms, and then be stitched together by a super-algorithm to deliver the main business intent. There is also a risk of lost business control, as a high-level intent may affect many of the lower-level business rules and processes. This effect is considerably smaller when an intent is introduced at lower levels, but in those cases there will be less business gain. The uniqueness of an intent and its context means there is little opportunity for reuse or experience building.

Ericsson believes that a new BSS architecture style that incorporates AI-native properties – including data-centric, learning loop, intent and event-driven logic, business rule hierarchy, and support for strategic, tactical and operational levels – is a much more efficient way to integrate AI with BSS. We fully agree with the view that future BSS and AI will be inseparably linked and must mature together [4, 5].

Introducing intents to BSS

An enterprise is a hierarchical or line-of-command structure in which business rules at the top steer, align and control the activities and behaviors further down in the structure. Business rules also steer all behavior in BSS, in pursuit of the goal of creating and maintaining a successful business. The step between a business rule and a business intent is very small; just a small shift of perspective.

A business rule is static and states what to do in a given situation, while a business intent states the desired or optimal outcome of a given situation – that is, interpreting what the stakeholder's interest is and trying to deliver as close to it as possible.

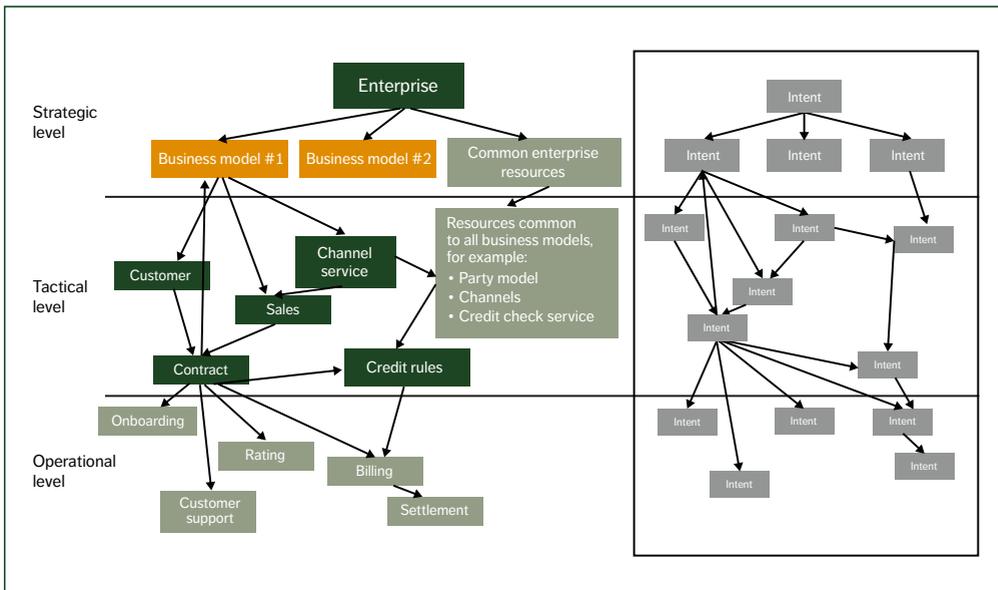


Figure 1 The business intent hierarchy mirrors the business rule hierarchy

Business intents can handle complex and dynamic situations and allow for feedback and comparison of actual and desired outcomes, enabling learning and knowledge building [1].

Higher-level intents in BSS are often expressed as business rules or KPIs. Intents are found at all levels of the business hierarchy, supporting both top- and bottom-line outcomes. An intent can range in complexity from an ‘atomic intent’ to an ‘algorithm of intents’ that combines a set of subintents. The term ‘atomic intent’ refers to the simplest possible intent structure, such as “our company will run a prepaid business model.” Note that an atomic intent on a strategic level

is likely to fan out into several intents on a lower level. While intents can be formulated for both human and machine consumption, they must be stated in a declarative format to facilitate automation in BSS.

Figure 1 illustrates the structure of the business intent hierarchy, which is designed to mirror the business rule hierarchy. An enterprise must have at least three different intent levels – strategic, tactical and operational [5] – all of which are the responsibility of the BSS. In software terms, these levels are equivalent to requirements, design and implementation, and execution.

	Strategic level	Tactical level	Operational level
Recurrence	One-off or few	Enough to learn	Very many
Human interaction	Yes	Yes	No
Learning/reasoning	Reasoning	Learning/reasoning	Execution
Feedback	Limited	Yes, key element	Large volumes
Main constraint	Quality	Quality	Time
Clarity of data use	Undetermined	Limited to own data	Deterministic and limited BSS set
Context, sources and data volumes	Many, external and large volumes	BSS	BSS and inference thereof
Data size	Large	Large but limited	Optimized for the intent
Data types and time series	All types	Many but limited	Few and limited to BSS origin

Figure 2 AI usage patterns are different at strategic, tactical and operational levels

The business strategy level owns and formulates the enterprise's top intents. For the sake of simplicity, Figure 1 breaks down only one business model to all three levels, but it is important to note that an average-size operator runs several different business models. The business tactic level is responsible for designing and implementing the intents of each of these business models at the operational level. This is normally achieved by breaking down the strategic intent into smaller, digestible parts, such as subintents with associated rules, information and processes. The business operation level – the core of traditional BSS – is responsible for executing to deliver the intents. This level requires further

automation to meet ownership and business transaction cost requirements. All three levels benefit from AI support but have different usage patterns and characteristics, as shown in *Figure 2*. The differences most worth noting are in terms of repetitions, context, exploration and data characteristics.

The OODA loop

The OODA loop [6] was initially developed in the 1970s as an in-combat decision tool of the U.S. Air Force. OODA stands for observe, orient, decide and act. Many of the OODA loop's basic concepts are found in today's software agent systems. The version

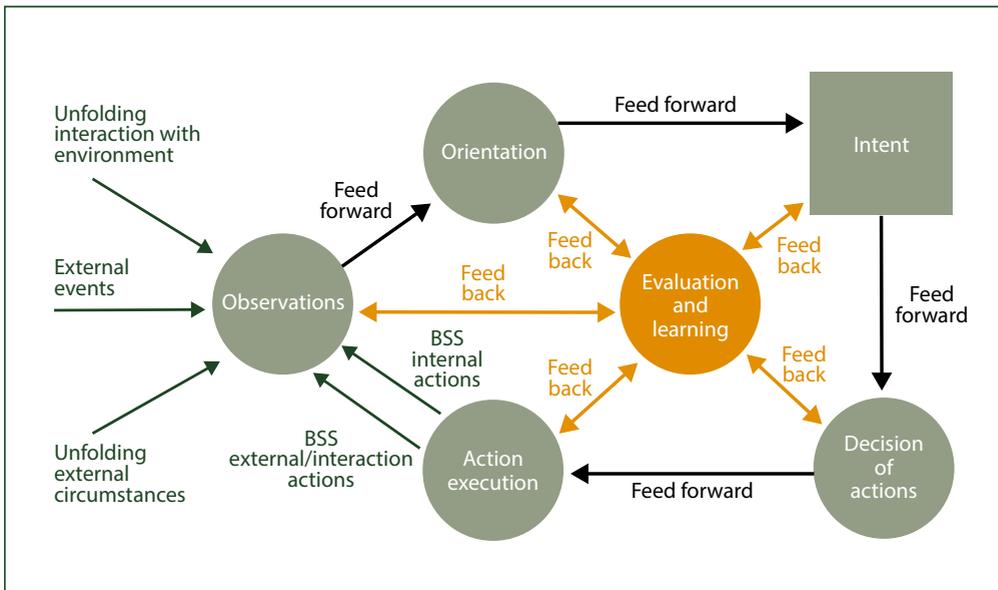


Figure 3 The OODA loop, modified to enable learning and intents

shown in *Figure 3* is complemented with explicit intent and learning capability.

The OODA loop idea is quite simple. First, observe or detect changes, events, stimuli or other things that happen in the context of interest, including internal states. Observations can be single or unfolding events, and they can be simple or complex in structure. Depending on the data, AI is often needed to interpret observations. Typical observations in BSS could be the availability of a customer's usage record or the arrival of a potential customer to a web shop.

The orientation step consists of aggregating and analyzing the observations that form the basis for the

decisions. Analyzing the individual observations and aggregating them into a complete situation description requires multilevel AI support. The orientation step in BSS should enrich the observation with customer data as much as possible. This data enables the BSS to select the correct rating and charging parameters for a particular customer when their usage record becomes available, for example, or to conclude that a visitor to a web shop is looking for a new phone but seems to be price sensitive.

To clearly separate the intent from the decision of action that fulfills the intent, we have added intent to our modified OODA loop. We achieved this by

dividing the traditional OODA decision step into two distinct process steps: intent and decision. Intents in BSS are statements of the desired business outcome in a given situation. In the case of invoicing, this would mean ensuring that the rates/charges on the invoice are in accordance with the customer's contract. In the case of a price-sensitive potential customer visiting a web shop to look for a low-priced phone, the intent would be to convince that individual to buy a phone in the medium-price range rather than choosing the least-expensive option.

The decision is limited by the inventory of available possible actions. A selection is made from the available array of actions that best matches the intent. It is also possible to enrich the decision with simulations to predict the action outcome. The action is simply the execution of the decision. In BSS, actions are carried out by business processes and they result in business outcomes.

Examples of BSS decisions (and resulting actions) would include the decision to apply the standard rate/charge process in the case of a new customer usage record, or the decision to show a price-sensitive web shop visitor not only low-priced phones but also medium-priced models that have received excellent customer ratings.

Evaluation, learning and feedback are essential to build a system with optimal performance that can adapt to both business and enterprise-external changes. An optimal system uses the process of orientation, intents and decisions to continuously compare and evaluate both business outcomes and its own capabilities. Adaption may require new or additional higher-level analysis, algorithm redesign and algorithm retraining. Sometimes, it is enough to have good in-operations learning, such as a continuous algorithm retraining capability. An example of learning in BSS could be reaching the conclusion that when dealing with price-conscious web shop visitors, better outcomes can be achieved by showing a mix of low-price and medium-price phones with good ratings, as opposed to including the expensive phones as well.

The OODA loop can have various depths of reasoning, from deterministic to deep learning – that

●● EVALUATION, LEARNING AND FEEDBACK ARE ESSENTIAL TO BUILD A SYSTEM WITH OPTIMAL PERFORMANCE ●●

is, the same OODA-loop engine can be used to observe, decide and select the proper actions for all event types, regardless of complexity. It can also be used recursively to build layered structures with arbitrary depth – that is, it can support multilayered business processes and interactions.

It is critical that the people working in AI-enabled processes are able to understand the reasoning behind AI-generated results. All of the steps in our modified OODA loop can be understood and executed by both humans and machines – which makes it possible to work together in the most efficient way possible based on the particular circumstances of the organization.

Use case: reducing manual handling of invoices

An invoicing-related use case provides a good illustration of how AI adds value to BSS. In this scenario, a telecom operator notices an increase in the number of invoices that require manual handling, which is costly for the company. The executive team initiates a strategic project to address the issue.

The objective at strategic level is twofold: to establish facts on the invoice situation and to state a future invoice strategy – that is, to set an intent – regarding invoice handling and cost. The strategy team pools information about known challenges in invoicing and gathers external data for benchmarking. With the help of classification and statistical analysis algorithms, the following strategic-level intents for invoices are established by the strategic project:

- » manual handling of less than 1 percent of all invoices
- » average handling cost per invoice of less than USD 1.

This use of two dimensions of intent ensures a sound business balance, helping to avoid potential pitfalls, such as the possibility of reaching 0.0001 percent of manually handled invoices at an average cost of USD 100 per invoice.

Once the strategic intent has been set, work on the tactical level can begin. The primary challenge at this stage is to understand and classify all the reasons why some invoices require manual handling while others don't, and to select the optimal AI algorithms that can deliver results in line with the strategic intents. The tactical level begins by defining an efficient subintent structure and estimating each subintent's yield to the strategic intent, in order to select the most valuable ones. Then, for each subintent, it classifies possible AI algorithms to identify the best ones. Important considerations include:

- » the information requirement and the complexity
- » the volatility of the constituent knowledge components in the problem, which in turn determines whether machine learning is enough or if it must be combined with machine reasoning or deep learning to create deep enough or adaptable algorithms
- » feasibility, effort and automation level in business operations
- » cost estimations, implementation, operation and support.

In this type of invoicing use case, it makes sense to introduce customized communication patterns that vary according to customer character type. Therefore, part of the work at the tactical level involves defining three distinct customer personas – angry, regular and docile complainers, for example – and customize anomaly invoice messages for each of them. The next step is to test these different messages on a small portion (1-3 percent) of the customer population to find the right message for each customer persona to ensure that their complaints are resolved without escalation to manual handling. Selecting only a small fraction of the total customer population reduces the business risk.

Finding the necessary knowledge components is an iterative task that requires AI support and access to relevant information. The tactical level is also responsible for the AI algorithm life cycle including design, implementation and operational launch. Further, it is responsible for stating the necessary changes to BSS, so it can both calculate according to the AI algorithms and automatically execute the new behavior in the invoice functionality.

The tactical-level subintents for the invoicing use case are:

- » invoice input correctness: higher than 99.999 percent
- » invoice anomaly statistics and predictions at both group and individual level: anomaly type, costs, volumes, services and dates
- » customer persona classification into three levels (angry, regular and docile complainers) with less than 1 percent error
- » customized message success rate above 90 percent.

The tactical-level changes to BSS in terms of new rules, information and processes are:

- » calculate invoice anomalies and recheck invoice input if anomaly probability is higher than 15 percent
- » determine customer persona complaint classification with continuous learning capability
- » customize the message success rate to achieve continuous learning capability
- » instruct customer support team to “cut it short” in cases with no anomaly and with angry
- » calculate and expose subintent and intent outcome, along with their projections and variance.

The methodology of the tactical level is similar to that of AI-supported CRISP-DM [7]. Once the tactical level steps have been completed, the role of the operational level is simply to execute the algorithms and the new BSS logic with as much automation as possible. After a short training period, the CSP can sort out the invoices that require

manual handling, identify the root causes, classify customers that systematically complain, and select the most efficient customized message. The error rate and the cost per invoice are initially quite high but decrease rapidly below the strategically stated intent as the invoice communication is tuned.

Implications

The invoicing use case makes it clear that BSS must have an omnipresent AI ability to be able to support strategic and tactical investigations, as well as having the agility in operations to change behavior to accommodate new algorithms, rules, information and processes. The inclusion of business models – that is, the grouping of rules, information and processes tuned to work together to deliver business outcomes for specific business situations – is also critical in the evolution of BSS. The intent structures must mirror the business model structures and their life cycles.

AI-native BSS require an expansion of the business logic elements – rules, information/objects and processes – to include intents and events. The business logic elements, often hidden inside applications, must be externalized to support the conversion of AI findings to automated BSS behavior. The business information must be structured in an ontology and made available to all business support users and applications, AI systems included, into a business information lake.

AI-native BSS must support both run-time and business-design-time. Traditional BSS are primarily

built with few configuration options for run-time, resulting in less agility [5]. While it is true that AI can help a business evolve in running BSS (for example in continuous development and operations mode), there is no avoiding the fact that this requires a BSS architecture that is at least partially new.

Conclusion

It is widely recognized that 5G and the IoT represent the main growth opportunities for communication service providers (CSPs) in the coming decade. To support emerging use cases in these areas, CSPs require business support systems (BSS) that can handle complex business situations and optimize outcomes with minimal manual intervention.

Artificial intelligence (AI) is the obvious answer, but introducing it into existing BSS is problematic for a number of reasons. Instead, Ericsson recommends an architectural change to traditional BSS to create AI-native BSS. Most significantly, this evolution requires the inclusion of an enterprise's strategic, tactical and operational levels in the BSS, together with the introduction of two new business logic elements (intents and events). One of the key differences between traditional BSS and AI-native BSS is the fact that AI-native BSS enable the various applications within the BSS to share business information with each other in an efficient and secure manner – a critical capability in the emerging 5G-IoT world.

Further reading

- » Ericsson, *The dawn of machine intelligence*, available at: <https://www.ericsson.com/en/news/2017/9/the-dawn-of-machine-intelligence>
- » Ericsson, *Zero-touch could herald a new era in service provider customer interaction*, available at: <https://www.ericsson.com/en/press-releases/2018/5/ericsson-zero-touch-could-herald-a-new-era-in-service-provider-customer-interaction>

The authors would like to thank Jörg Niemöller for his contributions to this article.

Lars Angelin

◆ is an expert in BSS within Business Area Digital Services at Ericsson. He has more than 30 years of experience in the areas of



concept development, architecture and strategies

within the telco and education industries. Angelin joined Ericsson in 1996 as a research engineer, and in 2003 he moved to a position as concept developer in the M2M and OSS/BSS areas. Since 2006 he has focused on BSS – specifically business support, enterprise architectures and the software architectures to implement BSS systems. He holds an M.Sc. in engineering physics, a Tech. Licentiate in tele-traffic theory from Lund Institute of

Technology in Sweden, and an honorary Ph.D. from Blekinge Institute of Technology in Sweden.



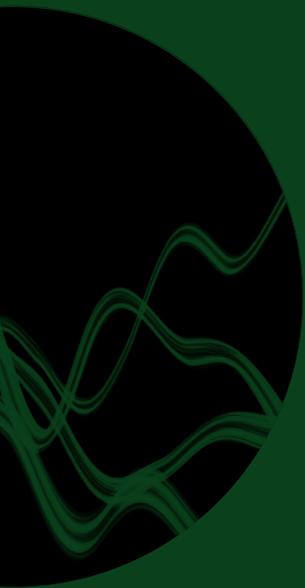
Johan Silvander

◆ is a senior specialist in information management who has worked at Ericsson

for more than 20 years. Over the years, his work has focused on the areas of OSS and BSS in a variety of different roles, including serving as a member of core architecture teams, working as a designer, taking technical responsibility for integration and installation projects, and being a test leader. He holds a Tech. Licentiate in computer science from Blekinge Institute of Technology, where he is currently pursuing a Ph.D.

References

1. TM Forum, **Open Digital Architecture, 2018**, available at: <https://www.tmforum.org/resources/whitepapers/open-digital-architecture/>
2. Ericsson, **Unlocking 5G's revenue potential: a roadmap for operators (press release), February 26, 2018**, available at: <https://www.ericsson.com/en/press-releases/2018/2/unlocking-5gs-revenue-potential-a-roadmap-for-operators>
3. Forbes, **How Artificial Intelligence Is Revolutionizing Business In 2017, September 10, 2017, Louis Columbus**, available at: <https://www.forbes.com/sites/louiscolombus/2017/09/10/how-artificial-intelligence-is-revolutionizing-business-in-2017/#58ebeb25463>
4. Harvard Business Review, **Artificial Intelligence for the Real World, January-February 2018, Thomas H. Davenport and Rajeev Ronanki**, available at: <https://hbr.org/2018/01/artificial-intelligence-for-the-real-world>
5. McKinsey, **Smarter analytics for banks, September 2018, Carlos Fernandez Naviera et al.**, available at: <https://www.mckinsey.com/industries/financial-services/our-insights/smarter-analytics-for-banks?cid=other-eml-alt-mjp-mck-oth-1810&hlkid=d3be9327efb84eccb44a1d8d391d0d8f&hctky=2669978&hdpid=ddc7fd1c-3815-4675-ae1a-f53e67d88452>
6. OODA loop definition available at: https://en.wikipedia.org/wiki/OODA_loop
7. CRISP-DM definition available at: https://en.wikipedia.org/wiki/Cross-industry_standard_process_for_data_mining



ISSN 0014-0171
284 23-3325 | Uen

© Ericsson AB 2019
Ericsson
SE-164 83 Stockholm, Sweden
Phone: +46 10 719 0000