WHITE PAPER

Enterprise Analytics
Data Center

Data In, Value Out

Identify and seize the opportunities to drive business value through advanced analytics innovation

Executive Summary

As new technologies like multi-cloud computing and the Internet of Things (IoT) emerge, businesses are generating more data than ever. This data is key to business transformation and enabling your organization to become a future-ready, insights-driven player. The key to extracting deeper, faster, more valuable insights from your data is analytics.

Today's analytics strategies must harness intelligence from increasingly distributed data environments, with important resources residing across the cloud, in your data center, and in connected devices sitting at the edge of your network. This means you should be making sure that you have a modern infrastructure in place that is optimized for analytics, scalable and able to grow over time. Implementing analytics is not a 'once and done' initiative – the type of analytics you need depends on business priorities and budgets. As your business evolves, so will your analytics environment.

This white paper explores five diverse business contexts in which analytics can play a role, exploring the technologies that enable them and considering real-world use case examples, as well as providing tips and learnings to help you explore further. Much of this is based on experience from Intel's own Information Technology group.

Table of Contents

About Intel IT

Intel IT is focused on using new technology to create business value and acts as a catalyst for organizational transformation.

As part of its commitment to driving industry-wide innovation, Intel IT has built its own competency center for advanced data analytics. This team is tasked with optimizing Intel's internal processes through machine learning and big data, targeting AI as the ultimate goal. The center then repurposes the technology to create products and solutions for use by its ecosystem and customers.

I. Critical Business Processes

1. Product Design

Advanced analytic techniques such as machine learning can help optimize the time-consuming validation stage and other aspects of product design by imitating and augmenting human validation capabilities. This helps to reduce the time taken to develop, test and launch your company's latest offering.

For example, to help enhance validation in its own computer chip design process, Intel IT developed a machine-learning platform called CLIFF, to uncover bugs in prototypes. The platform quickly browses through many thousands of historical test records to uncover patterns, a task that would take human reviewers thousands of hours, making it impractical to perform manually. Compared to standard regression tests, CLIFF validates the targeted functionalities sixty times more and identifies 30 percent more new issues on each run.

By automating the process, CLIFF has contributed greatly to Intel's strategic goals of reducing product validation time and reducing the number of iterations required. As a result, it has significantly shortened time-to-market and improved product quality.

The platform does all this using a form of prescriptive analytics. This means that it is able to not only predict likely outcomes through its machine learning algorithms, but also to inform and automate decisions about how best to tailor the process for each test moving forward. Even when petabytes of data are involved, the algorithms automatically identify historical patterns and take corrective action by adjusting the content of the validation test to increase the likelihood of finding a new bug in the hardware.

Looking Ahead:

The introduction of a prescriptive analytics tool like CLIFF is the first step towards shrinking the time needed for validation and so reducing time to market. The focus of this use case has been on relieving human testers of the burden of conducting repetitive but highly accuracy-critical tests through intelligent automation. With CLIFF, Intel IT's testers can work fast and efficiently, and re-focus their efforts on more value-adding tasks.

As technologies evolve, Intel expects that subsequent phases of this journey will create opportunities to not simply improve existing processes, but to augment human testers' abilities and enable them to drive more innovation, deliver more new products, and further reduce time to market. The phases we expect to see next include:

- **Intelligent Test Execution Management (ITEM):** Intel IT is now investigating ITEM as a companion initiative to CLIFF. CLIFF adds new tests to validation processes to ensure full test coverage, while ITEM analyzes the tests that are running, assessing which ones are actually uncovering bugs or other useful data. Any tests that ITEM finds are not adding value are removed, meaning that each validation process has full coverage while constantly adapting to eliminate inefficiencies.
- **Aided Debug:** Any debugging initiative requires you to analyze detailed bug data to understand the problem before it is possible to resolve it. The next step on from ITEM, Aided Debug will rely on machine learning to find the root cause of many of the issues, so that human specialists can focus their energy on developing creative solutions.
- **Smart Execution:** Once the stages described above have been reached, Intel's next objective will be to combine smart prescriptive analytics algorithms with the contextual information that only people can provide. Even the smartest algorithm will only be able to make decisions based on the data it can access within a system. People have a lot more peripheral context - for example, news about other initiatives going on at the company this week, or hearing from a colleague at the watercooler about a new change that has just been made in another area of the product. In the future, methods will be needed for workers to give this context to the system to enable it to further augment their own roles, creating a collaborative man-machine paradigm.

Advanced analytics helped Intel save USD100 million in unit testing and boost manufacturing quality.

2. Manufacturing and Quality Control

Quality control in a typical manufacturing environment often involves a number of different steps. Intel takes the quality of its products very seriously, and so has invested significant time and resources into ensuring each one of these stages is as thorough and efficient as possible. Analytics plays a key role, enabling more data to be processed faster and with greater accuracy than manual processes. For example, Intel IT has identified a way to reduce the overall number of tests without compromising quality. This has shaved seconds off this testing stage for each unit, which significantly speeds up the manufacturing process while reducing costs.

By replicating these sorts of process improvements across the whole manufacturing environment, even small improvements can add up to big savings and efficiency gains. This approach is referred to as personalized unit testing and works on the assumption that, like people, each unit is different. It requires its own combination of tests and decisions, so using analytics and machine learning to tailor the quality control process to each unit brings value comparable to that which personalized medicine brings to us.

Since introducing an in-memory machine learning platform to support these test stage enhancements, Intel has seen a reduction in the number of tests required, as well as improving the efficiency, reliability and scalability of the process. Intel saw a USD100 million saving in unit testing in 2016, as well as an improvement in manufacturing quality.

Looking Ahead:

Quality control measures in manufacturing typically still require a lot of manual intervention and human decision making. Experts must create content to be used in each test, research root causes of any issues they identify, and make decisions about what steps to take to drive improvements. A next step in the use of analytics in this process will be the application of machine learning algorithms to relieve some of this burden and automate some of these tasks so that

Figure 1: Business impact of advanced analytics solutions on Intel's sales and marketing.[1]

humans can focus on those aspects which machines cannot do. By introducing this technology to its own manufacturing environment, Intel expects to enable its employees to become more productive, and as a company to be able to test more products, faster.

3. Sales and Marketing

Sales and marketing teams need to transform vast quantities of data about their covered accounts into trusted insights, in real time. By using advanced analytics capabilities like machine learning, you can empower them to uncover fresh insights that were previously hidden within their data. This can help them have more effective conversations with their customers, discover and convert more leads, and enhance existing customers' loyalty.

Intel IT worked to create a machine learning-based market intelligence system and recommendation engine[2] that helps its sales and marketing division identify which resellers will connect most effectively with customers in specific industries and so enhance their ability to engage with and support those partners.

The original tool, called Sales and Marketing Account Recommendation Tool (SMART), provided the sales team with information about each reseller and its market, which products to offer them and opportunities to cross- or up-sell, based on its understanding of what has worked well with similar organizations.

A proof of concept (PoC) of the tool's second phase was recently completed in the EMEA online sales center. Building on the original tool, the latest version now combines Intel's CRM data with unstructured public data sources like news publications, patent filings, and information on hiring, venture capital funding, and merger and acquisitions. Leveraging AI technology such as text analytics, the revised tool can scan these extensive, disparate data sources and convert them into actionable insights for salespeople by imitating humans' language and reasoning capabilities. This volume of data would take months for humans to process, by which time a lot of it would be obsolete; by contrast, the new tool completes the analysis in just a few hours.

Intel has since deployed the system to cover its top 50,000 reseller customers in eight different languages. As a result, twice as many resellers in the engagement chain advanced from leads to qualified leads in comparison with the rest of the sales pipeline. These resellers also showed a three times higher click-through rate for email newsletters, and completed Intel training at a rate three times higher than the rest of the pipeline[3] - see figure 2.

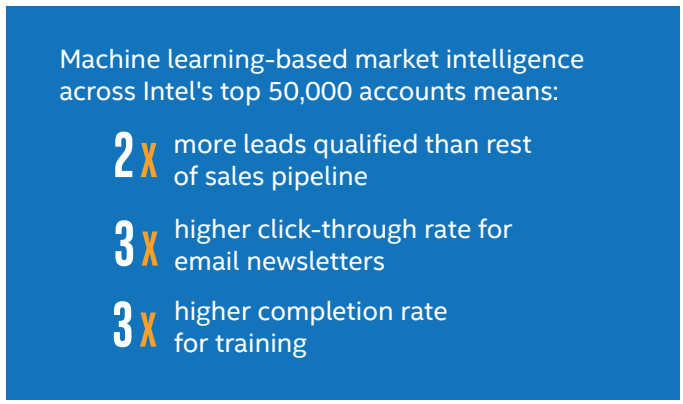

Figure 2: Effect of machine learning-based market intelligence.

Looking Ahead:

With the SMART tool in place, Intel aims to provide salespeople with a virtual personal assistant that will help them have more insightful and productive conversations with their customers, and deliver more value. The efficiency gains of having all this information at their fingertips will also enable salespeople to increase both the quantity and quality of their interactions.

Insights within a machine learning system like this flow in both directions. Each time a salesperson uses the tool, new data is delivered back into the system in return, enabling a constant process of algorithm fine-tuning. To this end, the Sales Assists tool is in a constant state of renewal, refinement and improvement. In addition, Intel IT is working to capitalize on these insights to build in email and other capabilities to help further improve dialogue with customers and increase opportunities for engagement.

II. Product and Service Innovation

1. Healthcare and Pharmaceuticals

With technology innovations such as wearable devices and advanced analytics, pharmaceutical companies can run more in-depth and accurate clinical trials faster and at a lower cost by enabling the continuous remote monitoring of patients. There is huge interest and potential in this area, with Ericsson* estimating that 4 million patients will use remote monitoring technologies by 2020[4]. Not only will this enable patients to benefit from new drugs becoming available sooner, but it is also expected to mean revenue per drug will jump considerably.

The Michael J Fox Foundation is using a wearable analytics platform, developed in collaboration with Intel[5], to help in its mission to find a cure for Parkinson's disease[6].

Each patient receives a wearable device equipped with accelerometer and gyroscope, which continuously captures data about the patient's movement. This data is securely transmitted to the cloud, where machine learning algorithms create objective measurements to accurately determine the drug's effectiveness. Using a series of Java* library tools, Intel has developed algorithms to track and analyze the data gathered by these devices, for example:

- **Activity level:** This algorithm measures the intensity of the device's movement, calculated as the average of absolute values of acceleration over 30-second intervals.
- **Tremor:** This recognizes and quantifies hand tremor by analyzing frequency, especially in amplitudes within the 4 to 12 Hertz range, and subtracts these typical tremor frequencies from the activity level measurement. A five-second segment with a high average difference between these two values is considered a tremor point.

Researchers, clinicians and data scientists can then use RESTful* APIs to extract, view and analyze the data. The Cloudera Enterprise Data Hub* is used to collect and manage the data, and Apache Spark* to manipulate the data before export. Some of the insights generated are also given back to the patient through a mobile app, which they can use to check their own activity levels and get updates or reminders about the medication and treatment plan - see figure 3.

By providing more objective, continuous information about patient symptoms during clinical trials of a new drug, the platform helps the Michael J Fox Foundation improve the quality and cost effectiveness of these trials. Teva Pharmaceuticals*[7] has also licensed the platform in a two-phase clinical trial for a new drug to treat Huntington's disease, a fatal neurodegenerative condition.

Looking Ahead:

The use of wearable technology is the beginning of an exciting journey in analytics and healthcare. Providing patients with always-on devices that are permanently connected to the cloud means pharmaceutical organizations can gain a vast amount more data than has previously been possible when relying on infrequent in-person check-ups with each patient. With more data, and the right machine learning algorithms and analytics processes, they can achieve results and deliver new drugs to market faster and cheaper. They are also able to use this ongoing data collection to proactively push updates, guidance and other communications to their patients, creating the effect of having a physician looking out for them all day every day, and encouraging patients to engage more actively with their own care.

In the future, opportunities to develop this personalized care approach will be huge, with other technologies such as video also having a role to play in both remote care delivery and deeper, more complex data gathering.

2. IoT Analytics

Smart manufacturing with IoT

Intel IT uses highly specialized IoT data specifications that make it possible to bring new factories online quickly and reliably by replicating identical processes from one factory to another[8] from implementing IoT devices to running complex analytics to automate and enhance operations, improve efficiency and reduce costs. In partnership with Intel's IoT Group and a range of industry-wide organizations, Intel IT

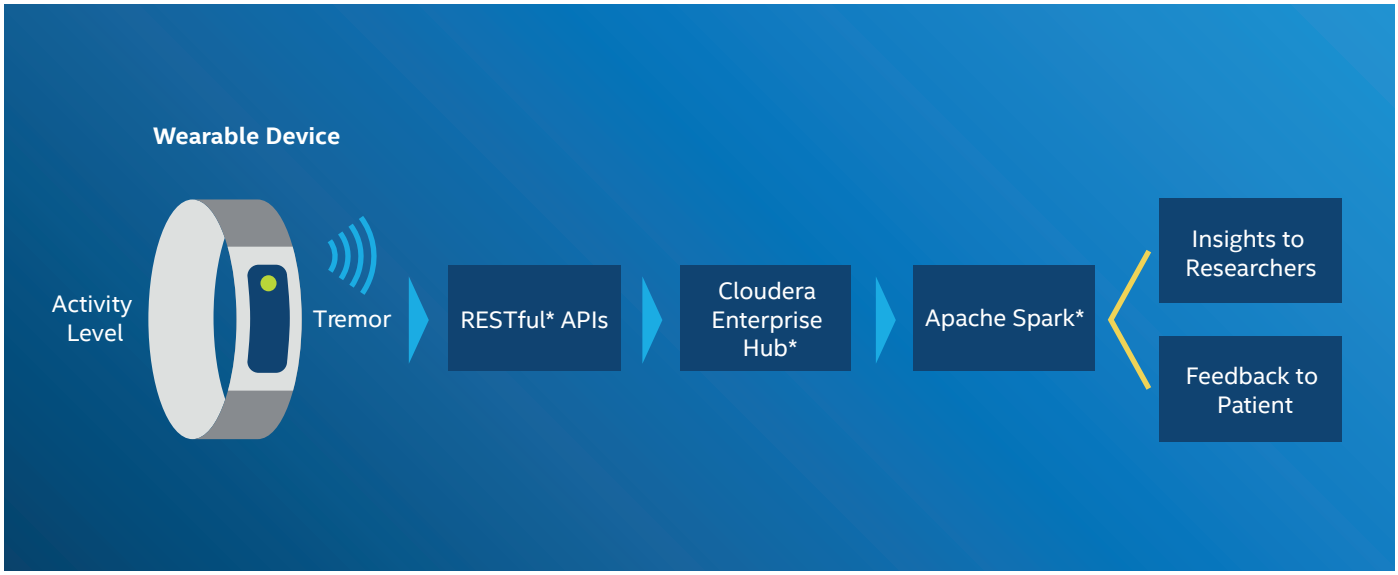

Figure 3: The wearable analytics platform process.

is establishing IoT data standards and messaging protocols for smart buildings that allow vendors to provide integrated solutions. With these standards in place, functionality in non-manufacturing building projects (which tend to rely on performance-based, rather than data-based measures) has the potential to become as reliable, repeatable and cost-effective as it is in Intel's smart factories. Analytics has a key role to play here, in turning the high-volume, high-velocity data generated through IoT into business-driving insights.

A blueprint for smart buildings

Intel IT is using IoT analytics to help it create a vision of the smart building of tomorrow at its Smart Building and Venue Experience Center in Chandler, Arizona[9]. The center acts as a testing ground for new IoT use cases as well as a working example of how IoT analytics can make a difference to business and operational efficiency today. Its aim is to develop IoT standards for smart buildings that will provide a blueprint for creating future smart offices and buildings.

All the building's systems - including HVAC, lighting, restrooms and the parking lot - are IoT-enabled (see figure 4). The center uses blueprints for various smart facility venues, developed by Intel IT and the Intel IoT Group, to create repeatable use cases and results which are then shared internally, with partners and with customers. For example, one solution monitors the parking lot to predict how busy it will be at any given time, helping facility managers control traffic flow and reduce congestion, while also enabling people looking for a space to find one more easily.

Looking Ahead:

An exciting area for further analytics innovation, IoT gives us the opportunity to make machines work for us. The patterns and insights found through predictive analytics and machine learning today will underpin the next step. As organizations evolve their analytics capabilities towards more prescriptive use cases, they will be able to automate many of the operational tasks in manufacturing, facilities management and a range of other areas that today take busy employees away from their more valuable core roles. This will not only help improve efficiency and cut costs, but it can also have a positive impact on the user experience.

For example, data from temperature and humidity sensors in a meeting room could be used to constantly optimize the environment – adjusting the temperature up or down, or even opening or closing windows – allowing those in the room to concentrate on their meeting. Or supply chain systems monitoring multiple internal and external data sources could identify where spikes in demand for a particular product are likely and automatically divert larger stock volumes to the affected area.

Figure 4: The smart building ecosystem

Predictive analytics has contributed USD 656 million in business value for Intel since 2012.

III. Learnings and Best Practices

Intel IT's own advanced analytics team has delivered dozens of projects over the last five years, and the incorporation of predictive analytics across the enterprise has contributed USD 656 million in business value[10]. In building and maintaining the advanced analytics competency center that has driven these initiatives, Intel has collected learnings and best practices that can provide invaluable foundations for future work:

Figure 5: Important roles in the data analytics team.

1. Build your team with distinct roles and shared capabilities

There are three clearly defined roles in data analytics teams: data scientists, product managers, and data engineers - see figure 5. At each stage of a project, each of these roles has a unique task. Teams must be structured around these three arms, with each member demonstrating a clear area of expertise.

It is also vital, however, for teams to share capabilities. This helps increase the speed and agility of the team as different skills are required at different stages of a project. Communication between team members is also much smoother when they understand each other's roles and skills.

In addition, having strong business acumen is essential to the success of a project. Even the most advanced technical competencies fail to add business value if you are unable to connect the dots and translate them into an improvement on the business outcome.

Focus on talent retention has helped Intel IT cut its voluntary turnover rate to just 5 percent vs a 13 percent industry average.

2. Focus on attracting and retaining talent

The science, technology, engineering and mathematics (STEM) skills gap has been well publicized, and corporate governance agendas frequently prioritize attracting and retaining talent. This is a competitive battle and it is vital for companies to remember that people are their most important assets. Often business priorities and investment behavior fail to recognize this, focusing solely on optimizing return on investment (ROI) and creating the fastest route to innovation.

At Intel, our people are at the heart of everything we do. We invest heavily in our teams, both financially and otherwise, to ensure not only that we take on the right talent, but that they are properly supported once they join the team. The Intel IT team has three people whose roles are dedicated to different aspects of HR, from full-time headhunting to nurturing team culture and development opportunities. This approach has helped us achieve a voluntary turnover rate of less than 5 percent each year – far lower than the US national average of just under 13 percent[11].

3. Stimulate a culture of innovation

It can be difficult for larger companies to stimulate the necessary culture of innovation for cutting-edge R&D and disruptive breakthroughs. The desire to maximize shareholder value can lead larger companies to adopt risk-averse business models. Start-ups, by contrast, are able to operate with speed and urgency to identify customer needs and deliver solutions. Their size allows them to adopt less hierarchical and more agile organizational structures, stimulated by incentives that reward collaboration and risk taking.

In order to adopt a start-up mentality and harness innovation in a larger company, it is vital to foster entrepreneurship. Intel IT has taken this approach. The team operates with a containerized, portfolio company structure and, although a centralized horizontal group, it is structured in six vertical sub-teams that each focus on critical and unique business practices – from product design through to marketing analytics. A lean HQ of only around 10 people sits above the separate verticals, adopting a function that is supportive rather than managerial, directed towards facilitating synergies between them.

This model has three main benefits: it affords a greater accountability for leadership within each sub-team by creating mini-CEO and supporting roles; it allows each vertical to operate independently, enabling them to adapt and evolve quickly; and it ensures each team is working at optimal efficiency with a dedicated, specific area of expertise.

Through this model, Intel IT has been able to adopt a very productive R&D mentality, developing breakthrough technology and scale at speed. We have also taken steps to eliminate risk aversion by making failure a measure of success: if the team isn't failing 30 percent of the time, then it's not being bold enough.

4. Demonstrate the value of your work: the 5/6/10 principle

Showing the positive business impact of your work is key to driving momentum around it. With the rise of analytics, many companies have made the mistake of spending millions on investing in infrastructure and tools to drive innovation without being able to show the tangible value of what they have created.

Figure 6: The 5/6/10 principle.

When Intel IT began to explore the ROI of analytic investments in 2009, the team decided to do the opposite. They chose not to invest in infrastructure and tools, but rather to focus their efforts on demonstrating value. They developed a '5/6/10 principle', taking five skilled people, giving them a problem to work on for six months, and in that time showing a USD 10 million ROI for the company - see figure 6. This approach allowed Intel IT to focus on low-hanging fruit rather than long-term goals, and therefore offer tangible proof of the value that machine learning could add to the enterprise.

5. Optimize your efficiency

a) Synergize internal- and external-facing activity

Product development projects generally focus on either internal or external use case potential. However, the most efficient approach is to think about both uses simultaneously. Amazon Web Services* is a great example of a hugely influential new service that began as an internal project, and – following its successful implementation and refinement – was then rolled out to external customers. Working in this synergistic manner has several benefits. Internal processes and appetite can be optimized through the exposure to competition that external-facing activity affords, incentivizing faster and more efficient workflow. And on the flipside, taking internal improvements to an external market can sharply save on R&D costs and offer the security of knowing that you have created a product that effectively addresses a need.

b) Be discerning about developing proprietary technology

It is similarly important to consider when to invest in developing your own technology, rather than taking advantage of tools already available in the community through resources such as open source. It would be foolish to waste money on buying a piece of software for millions of dollars if an equivalent tool is available for free. Aim to strike a balance that enables you to build reusable engines that can easily be adapted and repurposed for new projects further down the line. Focus your R&D investments on products and services that can't be found in the outside community – try to identify things that others don't see, and develop those into your USP.

c) Act when the time is right

Just because you have an idea, doesn't mean it is what the marketplace needs right now. Sometimes customers are not ready for the next phase of innovation, and it may be best to hold back. This could be on a level of demand (designing a new tool for which there is only very limited need), psychology (designing a tool that meets an existing need, but for which customers simply aren't prepared) or even practicality (tackling a problem that will drain resources faster than the returns it offers). Making the right decision on when to enter a new domain is vital to your business success.

6. Foster a relationship with industry and academia

An external perspective is vital to ensuring you remain one step ahead of the game. You might have the most impressive culture of innovation, but without staying attuned to activity in the wider world, and especially your competitors' movements, you risk going astray. A very effective way of ensuring this is fostering a strong relationship with both industry and academia. The team at Intel IT is encouraged regularly to attend industry conferences and other external events, and we are also heavily engaged with two top universities in Israel, contributing to masters' degree teaching and research in several different fields of data science. Fostering these relationships ensures our team stays well-connected and at the forefront of perception of the field.

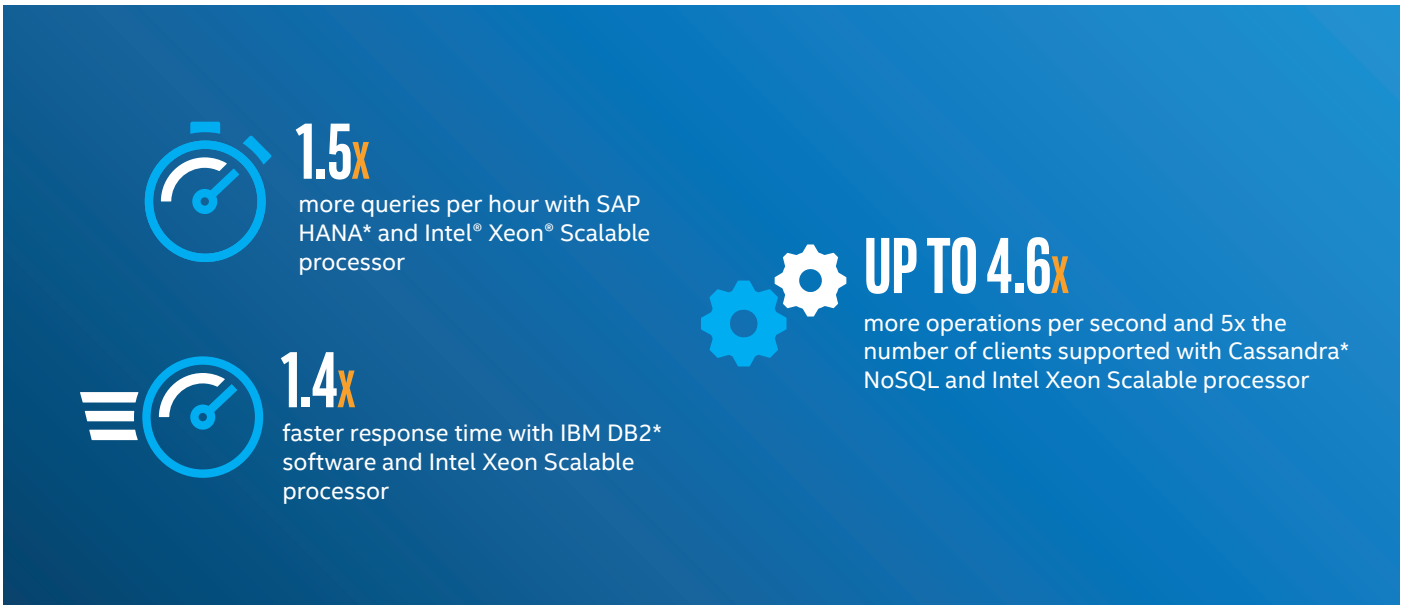

Figure 7: Accelerating analytics workloads with the Intel® Xeon® Scalable processors.

IV. Intel® Technology Enabling Analytics Innovation

In order to start developing your own advanced analytics initiatives, it is essential to have an infrastructure that delivers the performance, security and scalability to support business requirements today and enable you to respond to new opportunities and challenges tomorrow. Intel has developed a complementary suite of technologies to support analytics use cases, and works closely with a broad ecosystem to bring a range of analytics-optimized solutions to the marketplace.

High-performance, Workload-optimized Silicon

The latest Intel® Xeon® Scalable processors have been designed to accelerate analytics workloads from traditional models to real-time, in-memory computing ("scale up"), highly distributed workloads like Hadoop* ("scale out"), and even combinations of these models such as streaming analytics. For example - see figure 7:

- SAP HANA* running on the new Intel Xeon Scalable processor achieved 1.5 times more queries per hour than on the previous generation Intel Xeon processor[12], meaning faster time to insights for business-critical applications
- IBM* has reported similar gains on its DB2* software, with 1.4 times faster response time[13]– both when compared to previous generation Intel Xeon processors
- For scale-out analytics, enterprises deploying the Cassandra* NoSQL database, particularly to older installed-base servers, can expect to attain up to 4.6 times the number of operations per second and support five times the number of clients when using new Intel Xeon Scalable processors[14].

Intel Xeon Scalable processors improve data encryption speeds with negligible impact on overall performance, allowing IT to maintain fast analytics while protecting information. All the reliability, availability and serviceability (RAS) features of the Intel Xeon processor E7 family are also enabled on the new processor, including the expanded Intel® Run Sure Technology. The platform also provides strong flexibility and scalability by enabling IT teams to scale to two-sockets and eight-sockets and to take advantage of increased memory bandwidth for both in-memory and distributed solutions.

Performance can be further optimized by combining the latest CPUs with Intel® FPGAs - programmable accelerators designed to boost the performance of large-scale data and analytics systems. Intel FPGAs enable fast data processing by providing customized high-bandwidth, low-latency connections to network and storage systems. They also provide compression, data filtering and algorithmic acceleration. As they can be re-programmed, Intel FPGAs enable you to keep up with the latest algorithms and neural network topologies to ensure your machine learning and AI applications perform strongly for the long term.

As part of its ongoing commitment to delivering workload-optimized compute platforms, in 2016 Intel acquired Nervana, a solution stack designed for AI workloads. The Intel® Nervana™ Platform includes all the frameworks and libraries you need to design, develop and deploy state-of-the-art deep learning models. Training videos and data, and access to experienced specialists provide additional support to get you up and running quickly and with minimal fuss. Meanwhile, Intel® Saffron™ cognitive solutions provide a machine learning AI platform that mimics humans' ability to learn, remember and reason in real-time, helping to reveal critical, actionable insights faster.

Breakthrough Complementary Technologies

In addition to compute performance, Intel delivers analytics-optimized storage and networking technologies designed for latency-sensitive, memory-intensive applications. These include:

- **Intel® Optane® Solid-State Drives (SSDs):** Based on Intel's innovative 3D XPoint memory architecture, these offer unparalleled response times and endurance compared to NAND and traditional SATA storage. Intel Optane SSDs also serve as system memory, which helps

accelerate data applications and gain new insights from larger data sets, making them well suited to analytics workloads. These analytics-ready SSDs can be used in combination with Intel 3D NAND storage to help you build a tiered storage strategy that optimizes cost/performance across your data center.

- **Intel® Ethernet Adapters:** These offer one of the broadest 10GbE selections in the industry. Supporting 1/10/25/40GbE Ethernet port speeds, they are backwards compatible, providing an easy path for migration to higher speeds when more bandwidth is needed, for example as analytics workloads scale up.

Enhanced Security and Reliability

A rich suite of technologies for data protection and encryption help boost security and ensure strong reliability for your data-rich advanced analytics workloads. For example:

- **Intel® Advanced Encryption Standard New Instructions (Intel®AES-NI):** This instruction accelerates data encryption on Intel® technology-based processors, giving your IT environment faster, more affordable data protection.
- **Intel® Trusted Execution Technology (Intel® TXT):** This scalable architecture provides hardware-based security technologies to help ensure more secure platforms; greater application, data, or virtual machine (VM) isolation; and improved security or compliance audit capabilities.
- **Intel® Platform Trust Technology (Intel® PTT):** This platform functionality for credential storage and key management is used by Microsoft Windows 8* and Windows 10*. Intel PTT supports BitLocker* for hard drive encryption and supports all Microsoft requirements for firmware Trusted Platform Module (fTPM) 2.0.
- **Intel® Security Controller:** Enables automated and dynamic security provisioning, policy synchronization, protection, and remediation for software-defined infrastructure (SDI) built on VMware NSX*. It is designed to seamlessly broker between cloud orchestrators, software-defined networking (SDN) controllers, Security Connected solutions from McAfee, and the applications that manage them.
- **Intel® Run Sure Technology:** This standardized set of features helps reduce the frequency and cost of server downtime and protect the integrity of your critical data.
- **Intel® QuickAssist Technology:** Improves performance and efficiency across the data center by offloading servers from handling compute-intensive operations such as bulk cryptography, public key cryptography, and compression.

Software and Ecosystem Support

Analytics software platforms - both open source and those from industry-leading independent software vendors (ISVs) – are optimized for Intel® architecture, helping you achieve optimum performance for your analytics workloads. The variety of choice creates the opportunity for you to select the right solution stack depending on the type of analytics you need, from traditional BI to advanced machine learning. Intel® architecture-optimized ISV solution stacks include those based on SAP HANA*, SAS, Oracle Exadata*, Microsoft and Cloudera.

We also contribute to leading open source analytics platforms and frameworks such as Apache Hadoop* and emerging platforms like Apache Spark*, Storm*, Flink*, Shark* and NoSQL*. This helps ensure that Intel platform innovations are proliferated across the open source ecosystem, providing an alternative approach to those seeking to minimize software platform lock-in.

Intel's developer-friendly software investments have helped optimize the performance of popular developer frameworks on Intel technology-based platforms, enabling you to support the needs of your enterprise developers for advanced analytics projects. For example, the Intel® Deep Learning Software Developer Kit (Intel® Deep Learning SDK) is designed to help get complex analytics projects up and running quickly. A number of Intel software libraries also support advanced machine learning workloads. These include Intel® Math Kernel Library (Intel® MKL), Intel® Data Analytics Acceleration Library (Intel®DAAL), and Intel® Python* Distribution. We have also carried out optimization for Tensorflow*, Neon*, Torch* and Theano*.

Where Next?

There are as many applications for advanced analytics as there are business problems to solve. Ensuring you have the tools and resources in place to enable your organization to develop and run the analytics it needs is critical to driving digital transformation and ensuring long-term success. With a consistent Intel technology-based architecture across your infrastructure, you can be sure of a predictable path to rapidly scale and evolve your analytics initiatives, a consistent software programming model for your developers, and a versatile hardware platform that can support diverse types of analytics and speed time to results. Building a modern, software-defined infrastructure on Intel architecture also provides the flexibility of running and seamlessly servicing a variety of analytics workloads across your enterprise on-premise or on hybrid clouds.

Read on to learn more about how Intel technology can help you create your future-ready analytics environment:

- Planning Guide: Getting Started with Advanced Analytics
- Future-Ready Analytics Business Brief
- Intel IT Annual Performance Report 2016-2017

For more information about Intel's data analytics solutions, and to review the most up to date content visit: intel.com/analytics

For more Intel IT best practices, visit: intel.com/IT

About the Author

Itay Yogev

Intel IT Director,

Advanced Analytics, and Founder of IT Advanced Analytics Competency Center,

Intel Corporation

As the IT Director of Advanced Analytics at Intel since 2009, Itay Yogev heads up Intel's IT Advanced Analytics team, which focuses on solving high-value business problems and helping Intel grow software services through differentiated analytics. Itay has over 15 years of experience in leadership positions in the domain of analytics, as well as experience and expertise in software engineering, decision science, and business acumen.

1 Referenced from the Intel IT Annual Performance Report 2016-2017, www.intel.com/ITAnnualReport

2 Data mining uses machine learning to rediscover Intel's customers, https://www.intel.com/content/www/us/en/it-management/intel-it-best-practices/data-mining-using-machine-learning-to-rediscover-customers-paper.html

3 Data mining uses machine learning to rediscover Intel's customers, https://www.intel.com/content/www/us/en/it-management/intel-it-best-practices/data-mining-using-machine-learning-to-rediscover-customers-paper.html

4 The Growing Availability of Wearable Devices: A Perspective on Current Applications in Clinical Trials, http://www.appliedclinicaltrialsonline.com/growing-availability-wearable-devices-perspective-current-applications-clinical-trials

5 The Michael J. Fox Foundation Accelerates Research to Cure Parkinson's with Intel and AWS, https://aws.amazon.com/blogs/publicsector/the-michael-j-fox-foundation-accelerates-research-to-cure-parkinsons-with-intel-and-aws/

6 Intel and MJFF work to find a cure for PD, https://www.intel.com/content/dam/www/public/us/en/documents/white-papers/using-wearable-technology-mjff.pdf

7 IT at Intel Annual Performance Report 2017 www.intel.com/ITAnnualReport

8 Improving Manufacturing with Advanced Data Analytics, http://www.intel.com/content/www/us/en/it-management/intel-it-best-practices/improving-manufacturing-with-advanced-data-analytics-paper.html

9 IoT Data Standards Provide the Foundation for Smart Buildings, https://www.intel.com/content/www/us/en/it-management/intel-it-best-practices/iot-data-standards-provide-the-foundation-for-smart-buildings-paper.html

10 Intel IT Annual Performance Report 2016-2017, www.intel.com/ITAnnualReport

11 Compensation Force: 2016 Turnover Rates by Industry, http://www.compensationforce.com/2017/04/2016-turnover-rates-by-industry.html

12 1.5x claim based on SAP HANA* internal S-OLTP workload: Up to 1.59x higher OLTP performance (SAP HANA internal S-OLTP workload) vs. Intel® Xeon® processor E7 v4. Baseline config: 1-Node, 4S Intel® Xeon® Processor E7-8890 v4 on Brickland-EX-based platform with 1024 GB Total Memory on SLES12SP1 vs. estimates based on internal testing on 1-Node, 4S Intel® Xeon® Platinum 8180 Processor

13 1.4x claim based on IBM* BDInsights (3TB, 12 users, intermediate & complex queries). Testing conducted on IBM DB2* software comparing Intel® Xeon® Platinum 8180 processor to 4S Intel® Xeon® Processor E7-4890 v2 and E7-8890 v4. Testing performed by Intel® and IBM®, April/May 2017. BASELINE: 4S Intel® Xeon® processor E7-4890 v2, 2.8GHz, 15 cores, turbo on, HT on, BIOS 38.R02, 1.5TB total memory, 96 slots / 16GB / 1066 MT/s / DDR3 DIMM, IBM XIV Storage (132 10K RPM disks), SuSE Enterprise Linux* 11.3 kernel 3.0.101-0.47.71. NEXT GEN: 4S Intel® Xeon® processor E7-8890 v4, 2.2GHz, 24 cores, turbo on, HT on, BIOS 335.R00, 1.5TB total memory, 96 slots / 16GB / 1600 MT/s / DDR4 LRDIMM, 1 x 800GB, Intel® SSD DC S3700, Red Hat En¬terprise Linux* 7.3 kernel 3.10.0-514.16.1.el7.x86_64. NEW: Intel® Xeon® Platinum 8180 processor, 2.5GHz, 28 cores, turbo on, HT on, BIOS 119.R05, 1.5TB total memory, 48 slots / 32GB / 2677 MT/s / DDR4 LRDIMM, 1 x 800GB, Intel® SSD DC S3700, Red Hat Enterprise Linux* 7.3 kernel 3.10.0-514.16.1.el7.x86_64.

14 Up to 4.6x more OPS and up to 5x more clients based on Cassandra* Stress Test. Comparing 1-Node, 2 x Intel® Xeon® Processor E5-2697 v2 on Romley-EP with 128 GB Total Memory on Cen¬tOS7.3 kernel 4.10.1.x86_64 using Apache-Cassandra-3.10, Oracle JDK1.8.0_121 (1.7TB compressed dataset, LZ4 compression) vs. 1-Node, 2 x Intel® Xeon® Platinum 8180 Processor on Wolf Pass with 192 GB Total Memory on CentOS7.3 kernel 4.10.1.x86_64 using Apache-Cassandra-3.10, Oracle JDK1.8.0_121 (1.7TB compressed dataset, LZ4 compression)