Anomaly Detection for SAP Landscapes

Deep-dive into your SAP operations using machine learning ensembles with Fujitsu' s SystemInspection Service for SAP Solutions

FUJITSU

intel

mrcc.ovgu.de

MRCC VLBA

AG WI SAP UCC VLBA Lab

Fujitsu's technical competence team for SAP and the Otto von Guericke University Magdeburg (OVGU) jointly bridge data science and industry innovations. A newly released anomaly detection approach couples' machine-learning techniques with domain expertise in order to create best-in-class customer experience.

Management Summary

Anomaly detection for SAP landscapes is now available as part of the SystemInspection Service for SAP Solutions standard offering. The mechanism allows identifying periods of irregular SAP operations with a special focus on performance. As the underlying root causes entail the potential to affect business continuity, it is extremely important to detect and investigate anomalies as early and thoroughly as possible. The developed mechanism leverages a machine learning ensemble approach that results from joint research activities of Fujitsu and the Otto von Guericke University Magdeburg. Successful evaluation, as demonstrated in this white paper, resulted in general availability to support root cause analysis for SystemInspection customers. The SystemInspection Service for SAP Solutions is a packaged offering by Fujitsu, which aims at measuring and analyzing SAP landscapes by combining data mining techniques with handcrafted consulting expertise. It addresses SAP infrastructure related use cases in the field of transformations, assessments, or trend analyses in a neutral and data-driven manner.

Introduction – Bridging Science and Industry

Fujitsu continuously seeks innovative technology and exceptional customer experience. In the domain of SAP, we have just been awarded with the well-acknowledged SAP pinnacle award, recognizing Fujitsu as an innovation driver around data management and analytics solution while ensuring high-quality customer experience [1]. The new release of our Fujitsu SystemInspection Service for SAP Solutions is a perfect example to demonstrate these capabilities. It combines Fujitsu's longstanding domain expertise with state-of-the-art research artifacts, contributed by a specially founded R&D lab inside one of Germany`s leading academic institutes in the area of SAP related research.

For the last 20 years, the Magdeburg Research and Competence Cluster (MRCC), which is part of the Otto von Guericke University Magdeburg, delivers best-in-class application and knowledge transfer services around the complete SAP portfolio [2]. For more than 7 years, Fujitsu is a major partner of the MRCC research family, having supported the publication of 15+ scientific papers at international conferences and journals.

This white paper outlines the most recent innovation that resulted from our joint research activities. It is backed by state-of-the-art machine learning algorithms which our scientists have ensembled to a unique automated anomaly detection methodology. The approach is capable of scanning thousands of measured data points in order to quickly bring operational anomalies within our customer's SAP landscapes to light. As these points are likely to be overseen as part of classical monitoring strategies, they entail potentially severe problems in later lifecycle stages. These problems are to be detected and avoided by means of Fujitsu's SystemInspection Service for SAP Solutions. Consequently, the new technique has become an integral part of our performance root cause analyses as being demonstrated in this white paper. Furthermore, we outline some details about the underlying machine learning algorithms, their training and evaluation.

The Fujitsu SystemInspection Service for SAP Solutions

The Fujitsu SystemInspection Service for SAP Solutions is one of our offerings with a strong focus on SAP Technology such as SAP NetWeaver and SAP HANA. It offers an efficient analysis and comprehensive consultation package for existing infrastructure environments at a fixed price. The goal is to get a complete understanding of the current workload, performance and related resource consumption and distribution in order to provide measures and guidance to optimize your SAP landscapes according to the business strategy and requirements.

Value Proposition

The Fujitsu SystemInspection Service for SAP Solutions creates a holistic inventory of complete SAP landscapes. It collects and analyzes numerous performance metrics covering physical and logical layers such as infrastructure components, transactions, or dialog steps. The measured information serves as a comprehensive foundation to provide data-driven decision support for any upcoming transformation, migration, transition, or hardware refresh project. In short, the SystemInspection service supports to

- Avoid over-provisioning and unnecessary investments in IT equipment.
- Increase the quality of services by helping to eliminate or circumvent performance bottlenecks.
- Deliver individual and short-term results with minimum effort for the customer.
- Create insightful basis for capacity planning of upcoming new requirements to make strategic decisions.
- Give clear recommendations for optimizing the IT infrastructure for SAP landscapes.

The service works with hardware of any vendor, may it be on-premises or off-premises environments. Furthermore, it is fully GDPR compliant as only technical system statistics are processed and neither personal data nor business data are collected.

Service Design

The SystemInspection service is designed to combine an automated data collection and processing with subsequent handcrafted expert analysis. We understand that data and algorithm output always require to be analyzed and interpreted by domain experts in order to derive value and profound decision support. Therefore, the service comes with necessary consulting around it. On a high level, we go through the four phases, depicted in the Figure 1.

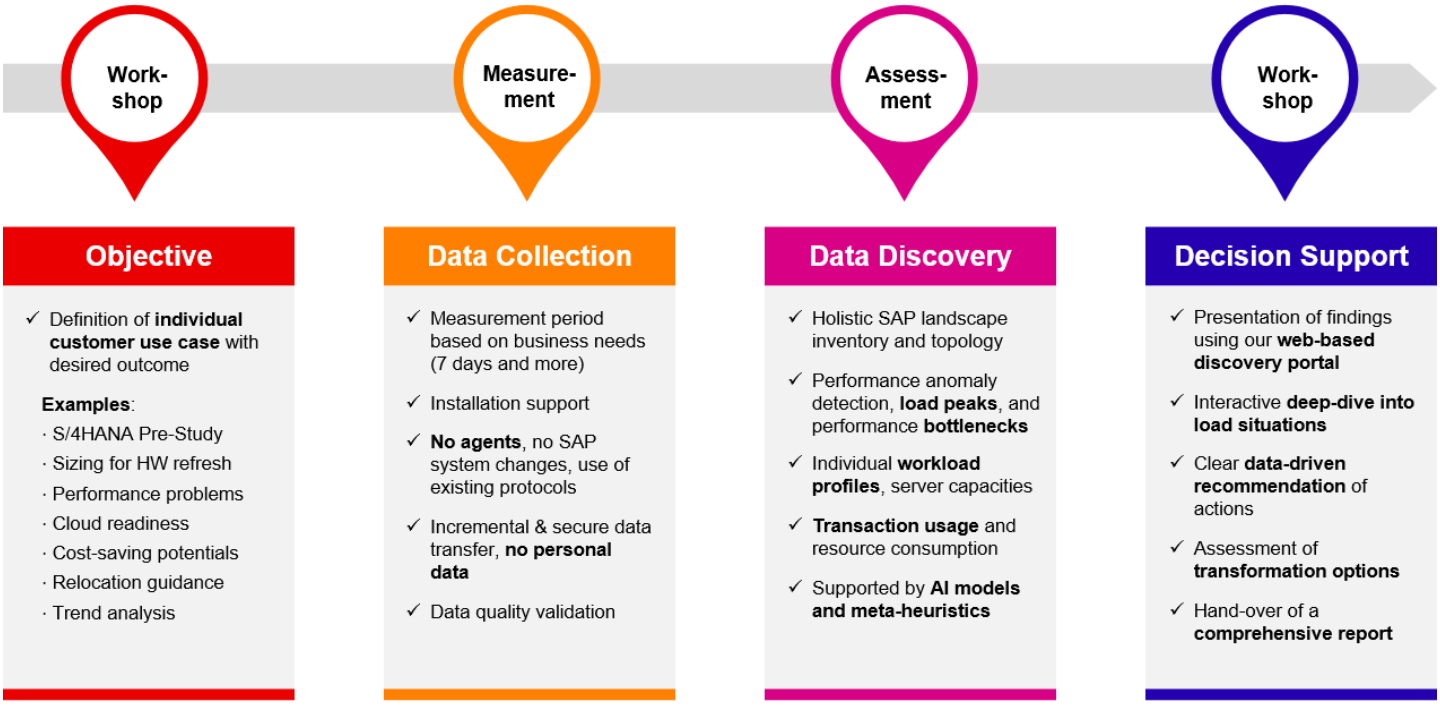

Figure 1: Fujitsu SystemInspection Service for SAP Solutions delivery process

The initial step aims at defining the individual **objective** our service is required to meet. Our experts will take the time to understand given challenges, strategies, planned transformations, or any existing performance problems. This way, we are able to align any upcoming activities with the customer-specific use case. Depending on the use case, a suitable period for the actual data measurement is to be defined. As an example, monthly or yearly settlements may need to be covered in order to get representable workload profiles. **Data collection** happens in the background without any noticeable performance impact. As our collector software uses existing standard protocols and stores all data locally, customers do not need to open their network or make security-related changes. As no master data or any other type of personal data is gathered, the service is also GDPR compliant. Once the data was uploaded to our analytics environment, we start the **data discovery**. This is where our unique research artifacts come into play. For example, the newly released anomaly detection technique, described in this whitepaper, helps to identify performance bottlenecks. However, algorithm output is always complemented by expert interpretation and individual preparation of the findings for a final workshop. Based on the initial objective, we jointly derive **decision support** from the analyzed data.

Data Mining

With the newly released anomaly detection technique, the Data Discovery phase of the Fujitsu SystemInspection service becomes even smarter. Together with the data scientists of our Fujitsu R&D lab, we went all the way through the Cross Industry Standard Process for Data Mining (CRISP-DM) [3].

Objective

The primary objective of our analysis is to simplify, automate, and improve the discovery of events in SAP operations that cause significant degradation in performance, which is noticeably hindering the efficiency of business processes. In other words, it is our challenge to identify anomalous points in time with respect to a number of performance metrics in order to prepare subsequent investigations of transactions and processes running at that time.

Data Collection

Data preparation requires to preprocess, and filter given data sets for a specific purpose. As the SystemInspection service creates a multi-dimensional inventory of a given SAP landscape, numerous metrics are available to feed learning algorithms for individual customer use cases. On a high level, the following categories are covered:

- Hardware configurations of the application and database servers such as SAPS [4] and memory capacity
- Software topology and software releases, kernel versions, patch levels, instance configurations
- SAP System load and resource utilization such as CPU and memory consumption, task type distribution, users, transaction throughput
- Performance metrics such as response times, dialog step quality, network throughput and I/O
- SAP HANA related metrics such as column and row store size, delta merge and savepoint performance, and detailed memory usage, e.g., for persistent memory [5] expert sizing

Metrics may be grouped by servers, SAP systems, SAP instances, SAP transactions, or OS processes. Furthermore, a time dimension allows to derive workload profiles for typical days or weeks of operation and to deep dive into specific points of interest. Figure 2 shows at a glance the type and volume of data which has been analyzed as part of Fujitsu's SystemInspection Service for SAP Solutions over the last years.

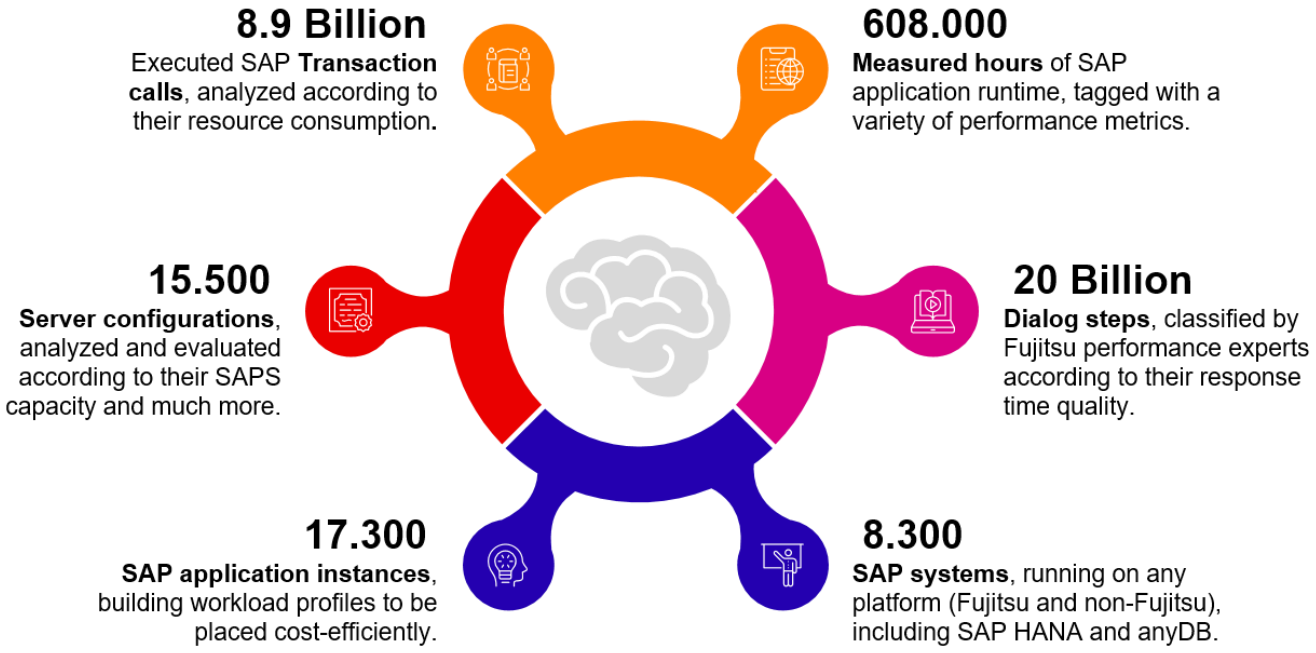

Figure 2: Data collected and analyzed as part of the Fujitsu SystemInspection Service for SAP Solutions

For anomaly detection in the domain of large enterprise applications, it is typically not sufficient to use only a single performance metric. For example, high response times may occur only rarely but do not necessarily indicate bottlenecks as they may simply result from higher complexity of the respective business transactions. If, on the other hand, performance degradations correlate with times of relatively low workload, the underlying components show unexpected behavior. Therefore, the present anomaly detection approach was designed to be fed with arbitrary metrics, depending on individual customer use cases. For the sake of demonstration, we focus on the following two metrics:

- Mean response time per dialog step in milliseconds
- Complexity classes of business transactions

Complexity classes represent a metric, which was designed by Fujitsu performance experts in order to rate the extent of load, caused by a single dialog step of an SAP transaction. By mapping complexity classes with response times, we derive another metric, which is termed dialog step quality. In the present use case, we make use of the ratio of good, moderate, and bad dialog step qualities.

As argued earlier, only common anomalies with respect to all of the chosen metrics are considered as performance anomalies within a given SAP system and period. In other words, we ignore rare and isolated events that might occur but ultimately do not have any notable effect on the business-relevant functions of the systems. Therefore, separate machine learning algorithms need to be executed before correlating their outcome. In the following, we outline some details about the so-called ensemble learning which backs our approach.

Data Preparation

As part of the data preparation phase, collected records must be transformed into a format that can be used in machine learning algorithms. In the first step, we convert the recorded historical performance data time series from continuous to discrete, with a time step of one hour. Therefore, all recorded measurements are averaged per hour. This aggregation effectively eliminates events that are too rare and which magnitude is not sufficient to notably affect the aggregated value. Additionally, discretization of the time series allows us to reduce the number of data points significantly. Such reduction of data points also dramatically reduces the computational complexity of the anomaly discovery. This, in turn, allows us to perform a faster analysis of the customer's data.

Our initial evaluation demonstrates that hourly discretization of the historical performance measurement time series is an adequate level of precision for our use case. At the same time, this discretization enables high performance and serves as a mean for eliminating the outliers, which are irrelevant for the specific business application. Furthermore, to account for the fact that we operate with relatively short periods of monitoring time, we use only weekdays and hours portions of time stamps on the recorded timeline. Specific months, years are discarded from training data as not providing any value for the specific application of machine learning. Additionally, many machine learning algorithms achieve the best performance or are even limited to computations involving only numerical or binary values. In order to account for this fact, we use a technique called one-hot encoding to transform day and hour values. One-hot encoding is a simple case of a concept borrowed from classical statistics, which is called Dummy Values [6]. Essentially, we treat weekdays and hours as categories, meaning there are seven categories for weekdays and 24 categories for hours of the day. Each of these categories is represented as a separate binary column in our dataset.

Modeling

The core of our anomaly detection mechanism relies on unsupervised machine learning. Unsupervised machine learning algorithms for anomaly detection, as opposed to supervised, do not require any prior knowledge on the specifics of the possible performance anomaly types and parameters. The obvious disadvantage of such an approach is that further qualified analysis is required to determine the severity and the root cause. When a performance anomaly is discovered, it is not automatically classified as a specific type. However, the advantage of unsupervised anomaly discovery is that even the anomalous system behavior will be discovered even if the specific performance issue was never encountered before.

While not requiring prior knowledge about the nature of anomalies themselves, most of the unsupervised anomaly detection algorithms still require expert knowledge for configuration. Often such configuration is done in multiple iterations until an optimal result is reached. This process is called hyperparameters tuning of a machine learning algorithm. In the specific case of anomaly detection, these hyperparameters ultimately control the sensitivity of the algorithm and how many data points should be considered anomalous by the algorithm. The overly tuned algorithm might not capture all anomalous data points, resulting in false negative anomaly detections. At the same time, an under-tuned might instead result in many false positive anomaly detections, which in turn lead to an increased time required to analyze results by the expert. In addition, the nature of most unsupervised anomaly detection algorithms is such that even when anomalies are not present in the dataset at all, if the algorithm is tuned to expect 0.05% of the dataset to be anomalous, even slightest deviations in collected measures will be falsely considered as anomalies to reach the expected amount.

To address this challenge, we have developed an auto-tuning mechanism for anomaly detection suitable for our specific use case. True anomalous events do not exist in isolation in SAP systems. When anomalous degradation in performance happens and affects business processes, the effects of this event are seen across multiple instances or can span multiple hours. In both

cases, a number of anomalous data points will be clustered together. Therefore, when we tune the algorithm such that it would detect more data points as anomalous, we expect the true positive clusters to grown. At the same time, false positive detections of anomalous data points will be unrelated to each other and dispersed across the time series.

The developed auto-tune mechanism starts with a set of parameters for the selected algorithm such that only a minimal fraction of the dataset would be detected as anomalous. Then, multiple iterations of the hyperparameter tuning are performed automatically. Each iteration is adjusting the algorithm's parameters such that more data points get classified as anomalous. In case of actual anomalous events existence, we expect more data points that are located next to each other on the timeline to be correctly classified as anomalies. This iterative tuning continues as long as clusters of anomalies are growing. If there are no more true anomalies to discover, then the data points newly classified as anomalies at the following iterations will be randomly dispersed across the timeline. If this happens, we stop the tuning and return the latest iteration results, where a growing cluster of anomalous data points was observed. If there are no true anomalies in the recorded workload performance, there will be no observed growing clusters of anomalies since the second iteration.

We have evaluated multiple anomaly detection algorithms for fitness in our use case. The most accurate results were currently achieved by employing a specialized anomaly detection algorithm called Isolation Forest (IF) [7]. The core idea of the algorithm is different from typical anomaly detection mechanisms, where algorithms attempt to model and predict the range of values that are expected in normal circumstances while classifying all that fall outside of that range as anomalies. IF attempts to isolate anomalous points and sequences while employing a decision-tree-based approach with an ensemble of these trees, hence the name Isolation Forest. The number of decision trees is a hyper-parameter defined before training the IF model.

Each tree is then operating with a portion of the whole dataset. Since it is a decision tree, it consists out of vertices. Each vertex is representing a value split. The specific value that is being used for a split is decided at random for every vertex. Specific value split is also decided at random.

Following visualization provides a simple example of this principle. There, we have a simplistic tree with five data points, where one data point is anomalous. Alternating values are response time and weekday.

During model training, a number of such trees are constructed. Anomalies are determined by the fewest number of required vertices to reach the specific data point. As mentioned earlier, internally the algorithm works by the principle of the ensemble of decision trees. Majority-vote of these trees needed to conclusively mark any data points as anomalous.

Such an approach demonstrates a high speed of model training and has moderate consumption of RAM. These advantages also enable us to efficiently run our custom hyperparameter tuning strategy efficiently.

Additionally, the sparse use of the computational resources enables us efficiently to use the supporting technical infrastructure of SystemInspection, ensuring the minimal impact on the overall system performance by the anomaly detection calculation.

Another notable challenge of classical anomaly detection approaches, which utilize unsupervised machine learning, is their potential lack of confidence in the achieved results, as different algorithms and hyperparameter combinations may return significantly different results. There are multiple strategies to addressing this challenge, often taking a form of an ensemble approach, such as combining multiple algorithms and averaging the results [8]. The goal of any employed strategy is to maximize the number of discovered true positive anomalies and minimize the number of discovered false positives. The best strategy is usually determined experimentally on the exact shape and the amount of data relevant to the specific use case.

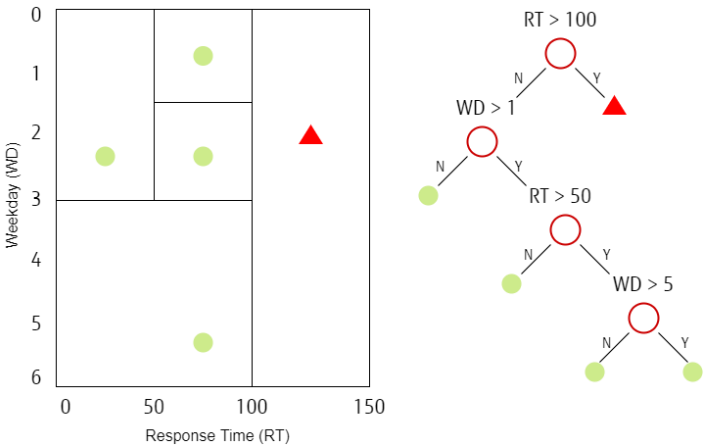

Figure 3: Internal data structure as used by Isolation Forests

We are tackling the challenge of confidence in a manner employing an ensemble of models. However, we are not using different algorithms. Instead, we train multiple models of the same selected machine learning algorithm, but each model is trained on a single selected performance metric or a subset of metrics. Then the results returned by all of these models are taken together and only time points where all of the models agreed on that there was an anomaly is counted as an anomaly.

It is worth noting that it is possible for a single model to be trained on all of the performance metrics simultaneously. Often that is a preferable option in applied machine learning. However, according to our experimental-based evaluation, we get more accurate results using the ensemble of models with subsequent deriving the intersection of the results of these models.

Additionally, depending on the algorithm, some hyper-parameters require retraining of the model while others can be adjusted in the existing model. This is important as training a model is an expensive operation. For example, in the case of IF, we have experimentally discovered that the only reasonably impactful hyperparameter for iterative adjustment is the contamination ratio. This hyperparameter can be adjusted in an already trained model. This parameter controls the threshold of the decision function used to determine which data points are considered outliers and which do not. This is advantageous for our use case as it enabled rapid calculation and iterative hyperparameter tuning.

Expert knowledge of SAP systems plays an important role in analyzing the achieved results of such an ensemble. It is not sufficient to just take all of the possible metrics from the recorded performance workload timeline and compare anomalies in all at the same time. Some of the metrics can be unrelated, and values have no meaningful correlation. To ease this task of analyzing the results, we pre-define the combination of different related between each other performance metrics. One example of such a combination is response time and bad dialog step ratio. When we take anomalies that were detected by the models trained on each of these individual metrics and only selected the intersection, we get meaningful anomalous time points.

Evaluation

Evaluation is a critical part of any research and development for applied machine learning in real world projects. We must be sure that the chosen algorithm and logic fulfill the business requirements. The best way to ensure that is to test it against a real-world baseline case where we know without a doubt an anomaly occurred. Comparing this baseline with the anomalies discovered by the selected machine learning algorithm, we can see how accurate the algorithm results are. We have used a real-world baseline where a known failure of one of the database nodes occurred during the SAP workload observation period. The exact time and effects on the overall landscape can be explored using SystemInspection standard functionality.

Two core performance metrics were selected for training. First is a transaction response time, representing how much time a specific transaction needed to execute. The second metric is a bad dialog step ratio. Anomalous values in each of these individual metrics do not necessarily signify an anomaly in the system performance. However, when considered together, the high ratio of bad dialog steps at the same time of spiking response times, signifies a noticeable, sudden anomalous degradation of performance.

Employing Isolation Forest in this specific case, the auto-tuning mechanism for hyper-parameters required four iterations of tuning the model until the optimal results were reached. During our evaluation, the most impactful hyper-parameter was the contamination ratio. Over the four training iterations, the following contamination ratios were evaluated: 0.01, 0.05, 0.1, 0.2. At the contamination ratio value of 0.1, the majority of the discovered anomalies were reasonably clustered along the time. However, already at the contamination ratio 0.2, a significant number of anomalies was discovered that were dispersed outside of the single anomalous cluster. These are false-positive anomaly classifications. Therefore, the tuning mechanism automatically stopped and reverted to the last known reasonable dispersion discovered, which was 0.1.

Since we train two independent models on two different performance metrics, we get two different sets of discovered anomalies. But, as expected, we observe an intersection of anomalies exactly at the time when a known system problem occurred in the baseline case. We have additionally executed evaluation anomaly discovery runs on a variety of simulated scenarios where anomalous events were injected into normal workload profiles. The results of these experiments reinforced the confidence about the high efficiency of the SystemInspection anomaly detection mechanism.

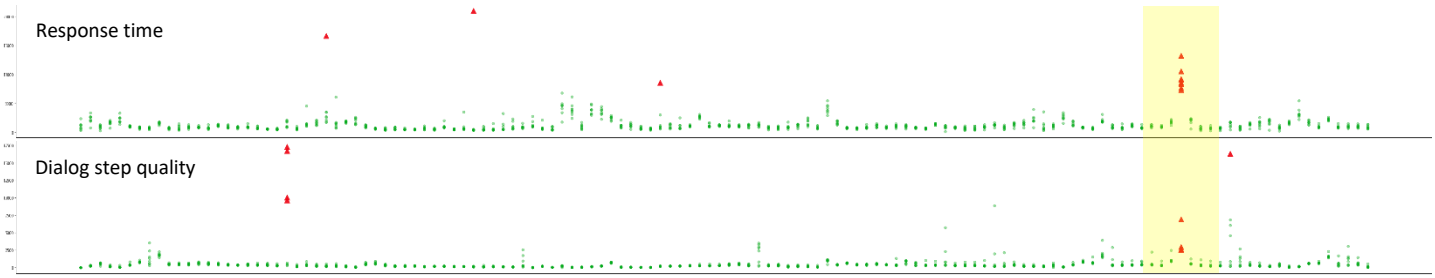

Figure 4: Interception of anomalies lead to points of business-critical operation

Decision Support

The SystemInspection Service for SAP Solutions provides clear recommendations and supports decisions in a data-driven way (see Figure 1, final stage). Therefore, the service surely does not end with the outcome of some machine learning based classification. In fact, results must be interpreted and further investigated by domain experts. To support this data discovery phase, our team of performance experts makes use of a web portal, which provides various means to analyze the measured data and deep dive into points or components of interest. On a high level, the following views are available to provide charts, metrics, and tables:

- Performance Dashboard: A set of high-level KPIs, which describe the overall health of the system landscape.
- Landscape Overview: Topology of the SAP landscape, versions, underlying physical or virtual servers, SAPS and more.
- SAP System View: Time-stamped data of performance, workload, and usage related metrics for each SAP system (SID)
- Host View: Time-stamped data of resource utilization and infrastructure related metrics for each server
- System Heat Map: Multi-dimensional view of selectable SAP workloads on servers to indicate landscape utilization.
- HANA Insight Analytics: SAP HANA related metrics: memory composition, delta merges, savepoint performance, and I/O.
- Anomaly Detection: Interactive 3D chart indicating anomalous time stamps for each SAP application instance.

The phases and views of any consultative investigation highly depend on the use case addressed by the service. A typical use case, which benefits from our newly released anomaly detection technique, is the root cause analysis of performance bottlenecks.

Root Cause Analysis

If performance problems are reported by end users, any additional information is valuable in order to find the root cause. For example, details on the transaction associated with the issues, the exact time and date, and information on preceding activities are beneficial. However, in many cases, problems are described rather generically, and user complains do not point to a specific event. Therefore, it is typically the initial challenge to identify any system activity, which has shown performance degradation compared to periods of normal operation. Secondly, these must be linked to exact timestamps in order to allow in-depth performance and workload analyses, reaching down to transactions, background jobs, or processes that co-existed at that time. The SystemInspection anomaly detection mechanism supports these steps by scanning all measured data points and separating the normal from the anomalous ones as explained in the previous section. Results are visualized in an interactive 3D chart (see Figure 5), which is part of the SystemInspection data discovery portal. This way, performance analysts are enabled to zoom into and explore the data points efficiently. In the example case, a productive SAP system, formed by eight application instances was measured for a period of three weeks. The anomaly detection mechanism revealed eight points of irregular performance at 3:00 pm, one for each application instance. Mean response times jump up to more than ten seconds as compared to normal response times below one second. This clearly indicates a system-wide issue. With the outputted information on the exact timestamp, further investigations are well targeted and the amount of data to analyze is reduced from more than 500 hours of operation to just one hour of interest.

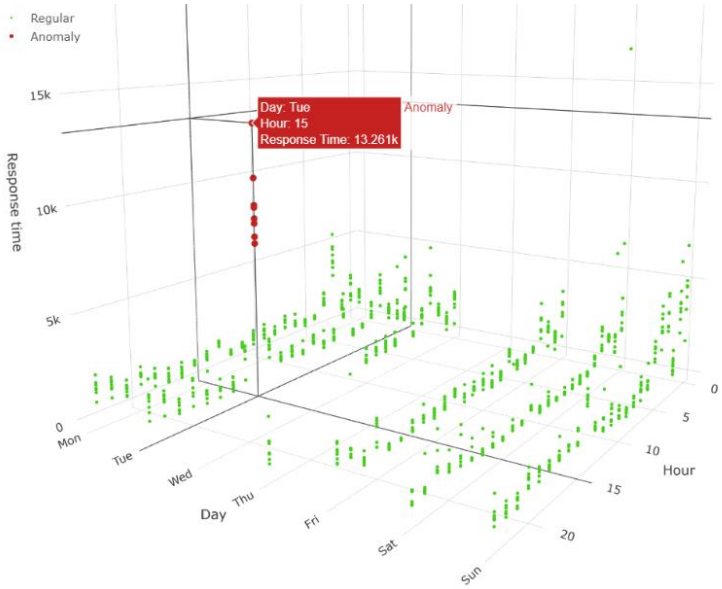

Figure 5: Anomaly detection interactive result in the data discovery portal

Subsequent steps of the root cause analysis involve several other features and charts, provided by the various views of the data discovery portal (see previous section). For the sake of demonstration, we cover only few of them in this white paper. One interactive way to explore performance metrics on a transactional level is a powerful, filterable and sortable table of so-called “Top Transactions”. It holds resource consumption metrics, grouped by any SAP transaction, executed within a measured system. It allows, e.g., to investigate the number of database accesses or the CPU share, taken by a transaction. In our given example, an extremely high portion of the overall response time was represented by database request times, indicating a potential root cause on the database layer.

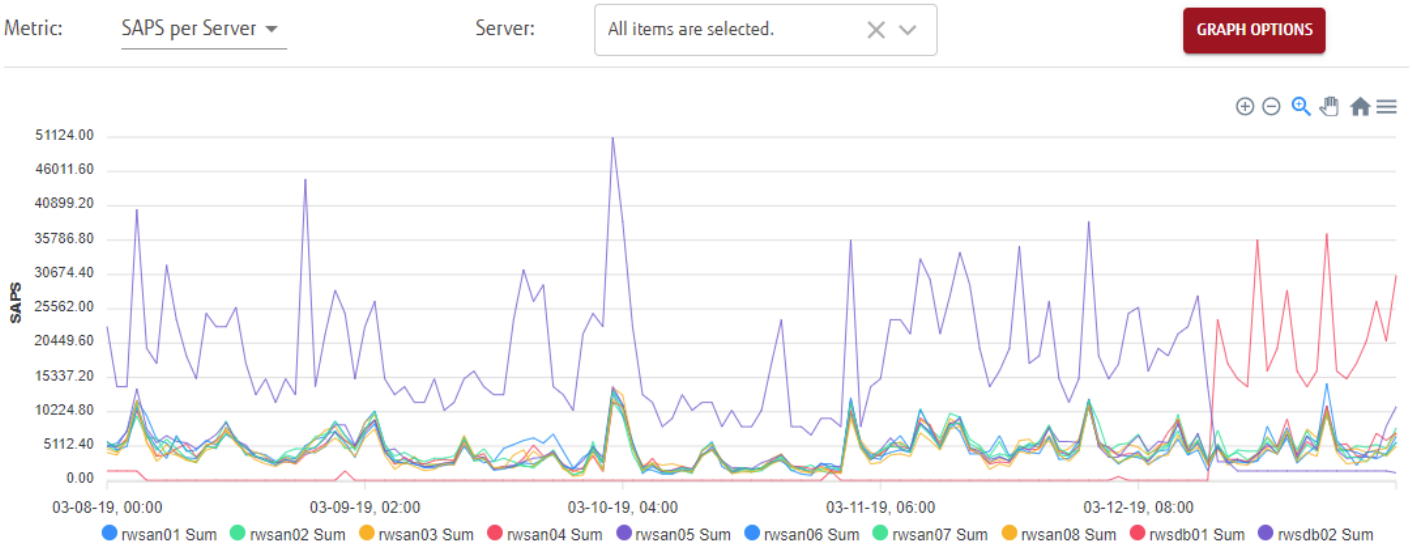

Figure 6: Workload profiles, grouped by server, indicating a database failover event

Following this path, the host view shows a number of metrics, which can be grouped by server. Similar to the Top Transactions, it allows to investigate any process that runs on the operating system (OS) layer. When mapping the timestamps of irregular data points (see Figure 5) with a list of running processes on the involved servers, the view reveals a suspicious distribution of CPU shares across OS processes. Particularly, a cluster synchronization process, which typically shows CPU consumption below 1%, has utilized the CPU by almost 30%. Since this faulty process heavily affected database request times, the customer was compelled to initiate a failover process on the database layer in order to resolve the issues. Such events are well represented by the measured workloads too.

For example, Figure 6 shows the metric *SAPS per Server*, which represents the CPU load (expressed in SAPS) on all eight application servers and the database cluster, formed by two additional servers. Putting a special focus on the database servers, we see load dropping on server *rwsdb02* while server *rwsdb01* seem to take over. Interestingly, this failover process follows exactly the period, which was classified by the anomaly detection mechanism to be irregular. To summarize, the SystemInspection service identified a faulty database process to be the underlying root cause for the performance issues that have been reported by application end users. Consequently, the customer's operations team was able to initiate proper measures with the clear objective to avoid a similar system behavior in the future. In this context, the customer reported the insights delivered by the SystemInspection service to be of very high value. In the demonstrated example, our customer was able to react to the performance issues quickly and consulted the SystemInspection team to identify the potential root cause. In other cases, existing minor bottlenecks may be unknown before effects become visible in a chain of faulty events. Therefore, it is extremely important to pay attention to anomalies as these may indicate bottlenecks or misconfigurations, which could lead to severe problems or, in the worst case, have the potential to interrupt business.

Summary of Deliverables

The SystemInspection Service for SAP Solutions is an analytics service, which addresses individual customer use cases. Therefore, any automated data mining features, as presented in this white paper, are combined with domain expertise. Recommendations, findings, and decision support is presented in a consultative workshop, by means of our data discovery portal. Furthermore, all metrics and findings are summarized in a comprehensive and structured report, handed to the customer and intended for reference and internal reporting.

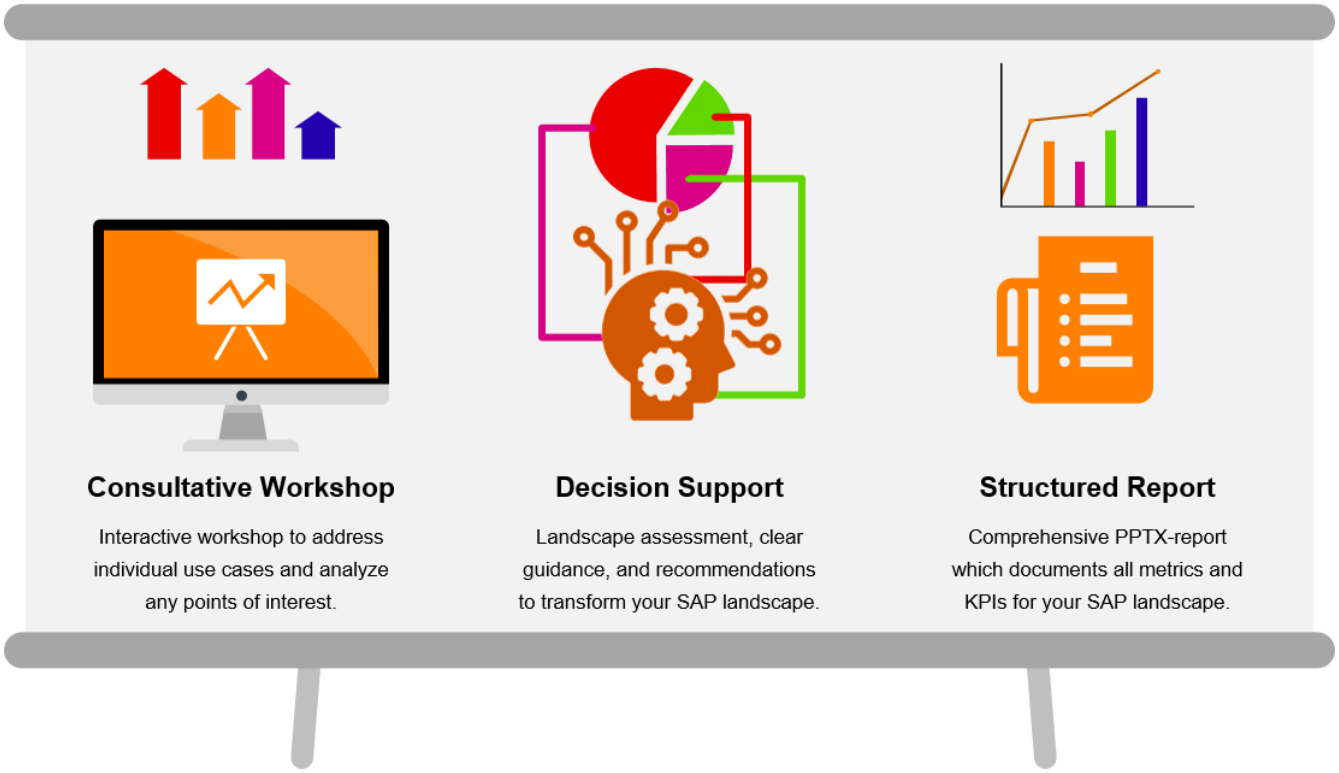

Figure 7: Deliverables of the Fujitsu SystemInspection Service for SAP Solutions

Acknowledgement

The authors of this white paper would like to thank Prof. Dr. Klaus Turowski of the Otto von Guericke University of Magdeburg for creating a research atmosphere that allows truly living the design science methodology and creating research artifacts that bridge science and industry. Similarly, at Fujitsu side, we thank Marcus Schneider to provide room and support for open-minded and research-driven innovations like this. Special thanks belong to Marco Edel, Jürgen Ellwanger, and Rohan D'Souza as continuous, trustworthy, and long-term supporters of the scientific cooperation between Fujitsu and the Otto von Guericke University of Magdeburg. Furthermore, we thank Intel for funding a Fujitsu PRIMERGY RX2540 M5 server with Intel Xeon Gold 6230 processors that ebabled us to train AI models in a matter of seconds. On a technical level, we also want to thank Christian Kowarschick who dedicated his outstanding SAP performance expertise to the research design phase. Furthermore, his continuous development efforts enable the SystemInspection team to efficiently collect and preprocess the required data from application, database, and infrastructure layers in order to deliver innovative data-driven consulting.

Authors

Dr. Hendrik Müller, Andrey Kharitonov, Jürgen Pfister, Michael Buchholz

References

[1] Award-winning SAP Partners. URL: https://www.sap.com/partner/find/award-winners.html. Last accessed: 05/2021

[2] Magdeburg Research and Competence Cluster, URL: http://mrcc.eu/mrcc/home/. Last accessed: 05/2021

[3] Chapman, P., Clinton, J., Kerber, R., Khabaza, T., Reinartz, T., Shearer, C., and Wirth, R. (2000). Crisp-dm 1.0 step-by-step data mining guide. resreport, The CRISP-DM consortium.

[4] Measuring in SAPS. URL: https://www.sap.com/about/benchmark/measuring.html. Last accessed: 05/2021

[5] Intel® Optane™ DC Persistent Memory and SAP HANA® Platform Configuration. Intel, SAP. URL: https://d.dam.sap.com/m/VuJWh4v/SAP%20HANA%20and%20Intel%20Optane%20Configuration%20Guide_final.pdf. Last accessed: 05/2021

[6] Suits, Daniel B. "Use of dummy variables in regression equations." Journal of the American Statistical Association 52.280 (1957): 548-551.

[7] Liu, Fei Tony, Kai Ming Ting, and Zhi-Hua Zhou. "Isolation forest." 2008 eighth ieee international conference on data mining. IEEE, 2008.

[8] İlker Kalaycı and Tuncay Ercan. Anomaly detection in wireless sensor networks data by using histogram-based outlier score method. In 2018 2nd International Symposium on Multidisciplinary Studies and Innovative Technologies (ISMSIT), 1–6. IEEE, 2018.

About Fujitsu SystemInspection Service for SAP Solutions

Anomaly detection for SAP is now available as part of the Fujitsu SystemInspection Service for SAP Solutions standard offering. The service can measure any physical and virtualized SAP environment, both on-premises and off-premises. The use cases focus on the following three focus areas:

- SAP landscape assessment: analysis of performance, bottlenecks, workloads
- Trend analyses and comparisons: before-after analysis of planned changes, growth trends
- IT Infrastructure transformations: requirement analysis for consolidation, migration, upgrade, transformation audits, expert sizing for Intel® Optane™ persistent memory

The SystemInspection Service for SAP Solutions leverages the power of data to address SAP infrastructure-related topics. To learn more, please contact cic@ts.fujitsu.com or visit:

https://marketing.global.fujitsu.com/sapsolutions

Contact
Fujitsu
cic@ts.fujitsu.com