Watch our predictions come true!

IT Operations Analytics (ITOA) – Optimize IT Operations with Big Data Analytics //

Timo Heise

Introduction

Companies are facing ever greater challenges. Globalization has triggered tough competition on an international scale, which puts enterprises under permanent pressure to change. Customers constantly demand new things, meaning that companies' products and services must be continuously adapted. To rise to these challenges and particularly in view of cost pressure, companies need to exploit new business models, and adapt and adjust their processes to cater to new situations. If companies' processes or business models change, then their IT landscape must be adapted accordingly. For IT operations, this means that the IT infrastructure must be highly available, scalable and be able to react easily to environmental changes.

Furthermore, end users should barely notice IT operations: Smooth and problem-free operations are expected to be the norm. For the chief information officer (CIO), this is a mammoth task, because current hot topics – such as the Internet of Things and the integration of cloud services – add to complexity, too. Despite these issues, the upper echelons of management continue to place high demands on their IT departments – they want increased efficiency, cost reduction, more digitization, and the fast provision of new and modern IT services. They expect a highly flexible and adaptive IT infrastructure to provide more agile processes while operating costs remain the same or even fall, thus enabling companies to react faster to customer demands.

The IT operations analytics (ITOA) concept addresses this issue and other challenges in IT operations. With the help of ITOA, it's possible to gain insights from different types of machine data (for example, log data) and thus predict future events. This way, data from IT operations can be used to improve service quality, increase the availability of the IT infrastructure, and forecast hardware component failures.

Why IT operations analytics (ITOA)?

Downtimes and malfunctions are expensive

If a company's IT infrastructure fails, all digital processes are affected, including connected production, which is why downtimes can be hugely expensive. According to market research company Gartner, an IT system failure costs an average of U.S.$5,600 a minute. The costs therefore escalate to over a million dollars in three hours. The study "Masters of Machines III – Mitigating the Impact of Critical IT Events" by the analysis company Quocirca also confirms Gartner's research. The study reveals that, on average, each critical downtime costs a German company €112,415. Such critical incident events (CIE) occur an average of three times a month and thus cause damages of around four million euros a year.[1]

System downtimes and malfunctions damage a company's reputation

An IT system failure or – in the worst case – destruction can cause a company's reputation to take a hefty blow. If customers can't use the services offered by a company, its image soon suffers. This is particularly true in the communications industry, retail, and the financial sector.

[1] https://www.it-daily.net/analysen/13658-kritische-it-ereignisse-millionenkosten-fuer-europaeische-unternehmen-studie

Error correction and root-cause analyses take time

If an error occurs in an IT system, it can be a while until the person who can fix it is identified. Tickets are sent from one department to the next and no one feels responsible. All this results in long error resolution times and dissatisfaction among users.

Performance problems

At certain times of the day, after advertising campaigns, or – for example – in the pre-Christmas season, IT incidents can cause poor performance due to the high number of users accessing the IT systems. In turn, this can lead to dissatisfied users or, in the worst case, lost revenue, for example, for online shops, because customers can't purchase the Christmas presents they wanted.

These are just a few reasons why companies should contemplate using ITOA systems. Fully linked IT machine data provides maximum transparency and simplifies and accelerates troubleshooting. This can then lead to cost reductions and shorter processing times for error messages. Another very positive aspect is that ITOA systems can predict imminent system failures and overloads by using modern forecasting methods. This is done by applying statistical models to the IT machine data and thus preempting and preventing failures. As a result, any performance losses can be counteracted and additional resources can be provided in good time.

What exactly is IT operations analytics (ITOA)?

With IT operations analytics, IT operating data is evaluated in real time with the help of big data analytics, so that the results contribute immediately and preemptively to improving IT operations. It is used in three main areas: **search, optimization, and forecasting**.

Every ITOA system has a **search machine** to scan all the data gathered. Usually, a specially developed query language is used for this. Rapid problem resolution and root-cause analyses are possible thanks to this search engine, which can perform evaluations across all the IT data.

Using modern machine learning methods, the ITOA system can learn the IT landscape's behavior and raise the alarm if there are any deviations. What's more, companies can deduce from the learned system behavior when further resources will need to be made available to avoid performance or capacity bottlenecks. ITOA can therefore also be deployed for performance **optimization**.

With **forecasting** methods, ITOA systems can calculate future system states and predict any possible failures. These can be answers to questions such as: When will a hard disk be full? Which hard disk will most likely fail based on historical analyses (analysis of smart data)? Will my system cope with a new advertising campaign in my online shop?

Furthermore, ITOA systems can be used to assign problems to specific IT operations teams and to prioritize them according to impact. In ITOA systems, alarms are also only triggered if system behavior deviates from the norm, and not based on fixed threshold values. This means employees are alerted only to really important incidents, and the system ignores brief fluctuations.

But which data is needed for which results?

The graphic below gives an overview of the possible data sources and the corresponding results and benefits that can be expected.

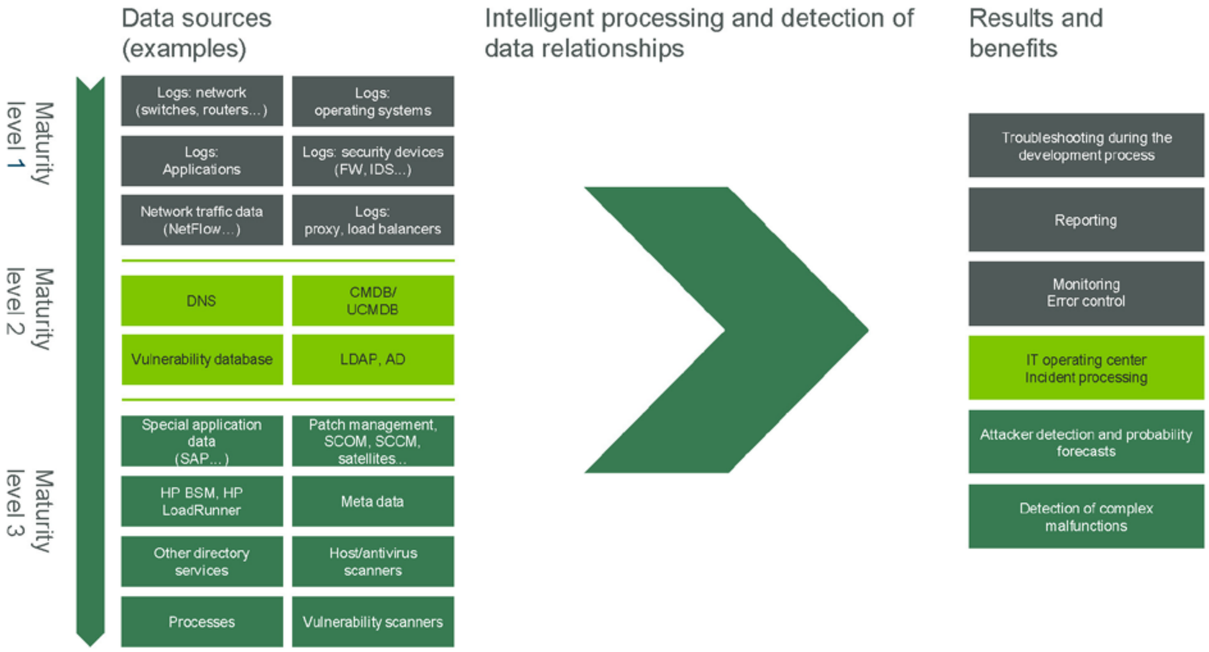
Data sources (examples)
Intelligent processing and detection of data relationships
Results and benefits
Maturity level 1
Maturity level 2
Maturity level 3
Logs: network (switches, routers...)
Logs: operating systems
Logs: Applications
Logs: security devices (FW, IDS...)
Network traffic data (NetFlow...)
Logs: proxy, load balancers
DNS
CMDB/ UCMDB
Vulnerability database
LDAP, AD
Special application data (SAP...)
Patch management, SCOM, SCCM, satellites...
HP BSM, HP LoadRunner
Meta data
Other directory services
Host/antivirus scanners
Processes
Vulnerability scanners
Troubleshooting during the development process
Reporting
Monitoring Error control
IT operating center Incident processing
Attacker detection and probability forecasts
Detection of complex malfunctions

In the first maturity level of an ITOA system implementation, log data is analyzed to simplify the search for errors during the development process, to establish reporting, and to enable cross-IT-system monitoring. The ITOA system acts here as a single point of truth: All the IT operating data is gathered centrally in one place and made analyzable. Thus, an end-to-end view of the IT landscape is achieved.

In the second stage of the implementation, more sources are hooked up, for example, an Active Directory or a vulnerability database. Once this stage has been completed, the ITOA system can be used as an IT control room and can help with incident processing.

In the last stage, application data – for example, from SAP systems – is integrated, and patch management and other directory services are implemented. This means the ITOA system can detect complex malfunctions and predict downtimes or incidents.

How and where is IT operations analytics used?

Identify problems 72 times faster with ITOA software (30 minutes rather than 36 hours)

A major corporation set itself the goal of ensuring very high availability (99.999%) for its own IT systems and identifying and rectifying any errors as fast as possible. A particular challenge here was that, although log data was collected, the different IT departments located across the globe used different analysis tools. The consequence of this was that, when serious issues arose, the employees responsible in the departments had to get together in “war rooms” to come up with a solution. And this meant up to 10 highly qualified employees were occupied for hours – and that was expensive. The company therefore decided to implement ITOA software to bring together the previous silos (13 different tools) and create a cross-department view of all data. Thanks to the ITOA tool, data could be analyzed centrally and across departments and errors could be identified and rectified earlier using automated log analyses, predictive analytics, and guided troubleshooting. Overall, error identification is 72 times faster with ITOA, resulting in fewer highly qualified resources being tied up.[2]

End-to-end transparency and fast troubleshooting with ITOA

A major telecommunications provider was faced with a challenge caused by the implementation of new middleware, which resulted in a very large number of tickets being generated in Remedy by the existing monitoring solution. The IT operations team was consequently unable to cope with prioritization and processing. There was also no end-to-end transparency and there were no KPIs for critical IT services until this point. The company therefore decided to implement an ITOA solution. The application and database log data, infrastructure metrics, network metrics, and the Remedy ticket system were then integrated with the ITOA system. By connecting these sources and creating statistical models, the ITOA system was able to prioritize the tickets automatically and assign them to the right IT operations teams. With the ITOA system, it was also possible to achieve end-to-end transparency and map the defined KPIs for critical services. Using the data analyses, the IT operations team could thus identify trends and discover anomalies, which could then be addressed proactively, preventing errors and downtimes.[3]

What do typical ITOA architectures look like?

An ITOA architecture must be highly scalable, have a large number of interfaces or connectors and a data store for historical analyses, and be able to perform evaluations in real time. As well as those requirements, security is very important. The graphic below shows a general, greatly simplified ITOA system architecture:

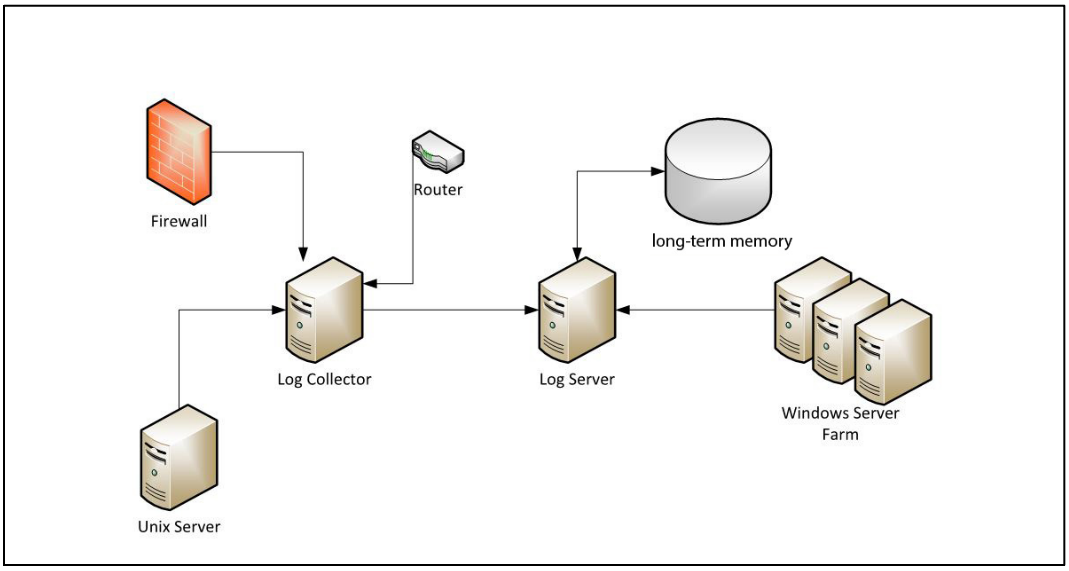

The log collector shown in the graphic collects all kinds of machine data in each network segment and transmits it to the log server. This data can be from firewalls, servers, switches, routers, or other machine data as shown above. The log server is responsible for access to the indexed data and the data's organization. In addition, a long-term store is connected to the log server so that historical data can be evaluated. The central instance here is the log server, which possesses data analysis capabilities. Depending on the setup

[2] https://www.hpe.com/h20195/V2/GetPDF.aspx/4AA5-4931ENW.pdf
[3] https://www.splunk.com/en_us/customers/success-stories/vodafone.html

and customer requirements, this solution can be extended with as many log collectors and log servers as needed, to ensure high performance and availability.

Which providers and products are there?

The ITOA software market is currently highly fragmented with around 90 products and there's no single "best" ITOA solution. Rather, it depends on the use case. Some providers – Nexthink, for example – specialize in end-user analytics. Others – such as ExtraHop – focus on the analysis of network data. There are also products such as HP Operations Analytics, IBM Operations Analytics, and Splunk, which enable a cross-company analysis of IT operating data and are already established and tried-and-tested on the market. To find the most suitable solution, we recommend setting up a detailed tool selection process with use cases specified in advance.

Some criteria are especially important when selecting a tool. A study by IDC identified key factors for choosing the right ITOA system[4].

Companies looking to invest in an ITOA system should keep the following criteria in mind when drawing up their shortlist:

Functionality

- The system must have an intuitive, modern, and well structured user interface.
- It must be possible to create dashboards.
- A function to set up alerts when thresholds are exceeded must be available.
- Historical data analysis must be possible.
- Log data searches must be possible.
- The system must support analytics functions such as machine learning or have interfaces to R or Python.

Security

- Audit-compliant storage of log data is required.
- Passwords must be encrypted.
- Communication between the source systems and the ITOA system must be encrypted.
- Communication between the ITOA system user and the ITOA system itself must be encrypted.
- A role-based authorization concept must be implementable.

Architecture

- Current platforms such as Windows, Linux, and Mac must be supported.
- It must be possible to import/connect log data sources.
- It must be possible to develop enhancements.
- The architecture must be highly scalable and enable big data analytics.
- The architecture must be able to perform evaluations in real time.
- The architecture must be protected against unauthorized access from within and from outside.

[4] IDC, HP Operations Analytics Simplifies and Accelerates IT Service Delivery and Digital Innovation, 2015, p.5, 8

The opportunities and risks of ITOA systems

Opportunities

An ITOA system can increase efficiency by allocating resources specifically. Furthermore, a centrally organized ITOA system can facilitate root-cause analysis, because searches don't have to be performed separately in the individual user departments, but can be executed in a central data store. Capacity bottlenecks, too, can be predicted using intelligent algorithms and then avoided through the provision of additional resources. Furthermore, one-off questions can be answered with ITOA systems, such as: "What problems were caused by a change in provider?" or "Which specific action (for example, a patch) triggered an error?" ITOA systems can therefore produce a holistic view of a company's IT environment and its current status, and therefore create transparency.

Risks

ITOA systems store large quantities of security-relevant data, such as access data. Furthermore, ITOA systems require a connection to each network participant and access to the associated data, which can pose a security risk if the system is not professionally configured. Security is therefore a crucial topic. It's very important to carefully protect ITOA systems from access by third parties, because once a system is taken over, other connected devices throughout the entire infrastructure can be compromised. A further risk is that the architecture could be overloaded due to the high frequency and quantity of data, preventing analyses from being performed. Similarly, high costs can play a role in this context, because the data volume grows incredibly fast and many ITOA software providers bill according to the data quantity indexed daily. Data protection is also paramount, because user activities at employee level could be measured using ITOA system data and the machine data collected can also contain personal data, such as IP addresses.

What happens in a typical ITOA project?

ITOA systems are complex, so it makes sense to use a procedure model when implementing one. The graphic below shows a possible approach and is subsequently explained.

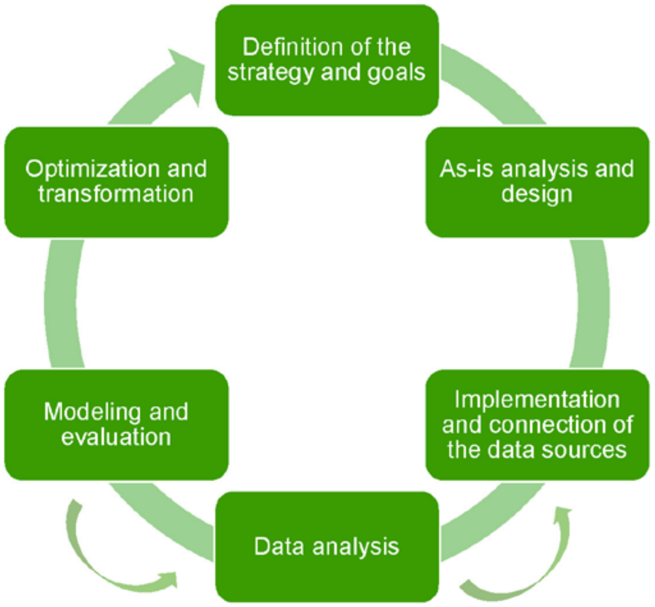
Definition of the strategy and goals
As-is analysis and design
Implementation and connection of the data sources
Data analysis
Modeling and evaluation
Optimization and transformation

Phase 1 – definition of the strategy and goals

In the first step of the project, the strategy and goals should be defined. To do this, the current issues and potential for improvements in IT operations should be gathered and an ITOA solution should be derived from these. The stakeholders in the ITOA project should also be identified at this stage. For example, an ITOA project team could comprise an ITOA architect, a system administrator on the customer side, an incident manager, a data protection officer, a data scientist, and a project manager. Not every role has to be involved in all of the project phases. This must be decided according to strategy requirements. A risk analysis also belongs in this stage.

The following tasks should be complete by the end of phase one.

- Goal of the project/definition of the problem statement
- Definition of the stakeholders/project participants
- Definition of the project risks
- A defined project plan

Phase 2 – as-is analysis and design

In the as-is analysis and design phase, the company's current infrastructure is evaluated and then, based on this, an architecture for the ITOA system is drafted. As well as definition of the architecture, an authorization concept is formulated here. Furthermore, an ITOA tool is selected at this stage. The previously listed minimum requirements can serve as a basis for this, and must be refined accordingly. As part of the design phase, the decision is made about which data sources should be connected to the ITOA system. To do this, the name, the data type, the location and owner, the format, the import method, the volume, the frequency and quality of the data, and the data usefulness must be determined. Once an architecture has been defined, a tool has been selected, and the data sources to be connected have been identified, the implementation phase can begin.

At the end of this phase, the following results must be achieved:

- Overview of the current infrastructure
- Design of an ITOA system architecture
- Definition of an authorization concept
- Identification and explanation of the relevant data sources

Phase 3 – implementation and connection of the data sources

This is where the architecture designed in the second phase is implemented and the source systems are connected. For optimal results when connecting the data sources, the maturity model should be used. The data sources for the individual maturity levels are not usually connected in one single iteration of the procedure model. Instead, it makes more sense to complete one maturity level after the other, to keep the complexity of the project to a minimum. We therefore recommend performing the ITOA procedure model in three iterations, to reach the third maturity level. However, if the objective of the project is solely to set up monitoring functionality, then one iteration is sufficient.

Every iteration of the model should include the cleansing and preparation of the data for an analysis. Here, the connected data sources must be prepared in such a way that they are in the right form for advanced statistical analyses or machine learning algorithms.

The following points should be completed by the end of this phase:

- Installation and configuration of the ITOA system in line with the architecture concept and the authorization concept
- Depending on the goals and the maturity level defined:
  - Connection of internal data sources (for example, logs from operating systems, switches, routers, and so on)
  - Connection of directory services and inventory databases
  - Connection of third-party systems (for example, application data, ticket system, monitoring system, and so on)
- Data preparation for analysis:
  - Data cleansing
  - Definition of data fields, if they are not identified by the tool

Phase 4 – data analysis

In this phase, the data sources are set up by defining queries in such a way that the current system status can be seen on reports and dashboards, and error monitoring is possible in real time. To obtain first insights, preconfigured apps or dashboards that offer provider-specific queries can be used at this stage. Phase four is where the first added value is generated for the company. Through monitoring, a current overview of all IT operations is achieved – and end-to-end monitoring is possible once there is a full implementation across the entire IT infrastructure. This phase is closely linked to the previous phase of data connection and preparation, so the company may realize at this stage that further data sources are required for a meaningful analysis. If this is the case, the project shifts back to the previous phase. The next phase – modeling and evaluation – also dovetails with this fourth phase, because the analyses created here form the basis for forecast models.

The following should be completed in this phase:

- Definition of search queries to answer questions from the strategy phase
- Identification and use of extensions and apps for rapid reports and an initial overview
- Integration of validated search queries into dashboards and reports, including visualization of real-time monitoring

Phase 5 – modeling and evaluation

This phase is about the use of predictive analytics. The first step is therefore to identify a suitable analysis tool and analysis methods. Ideally, an ITOA solution should contain this functionality. But if not, forecast models can be created by connecting or using common programming languages such as R or Python. After modeling, the models must be validated and the results evaluated. This phase is repeated any number of times depending on the question.

The results of this phase are:

- Validated statistical or data mining models
- Forecasts about possible capacity bottlenecks in the system, new relationships that cause malfunctions in the infrastructure
- Implemented models in the ITOA system that analyze the incoming data continuously during operations and make forecasts in real time

Phase 6 – optimization and transformation

The last phase is about using the insights gained for IT operations. Based on these insights, identified hardware bottlenecks can be fixed by purchasing additional hardware or troubleshooting processes can be improved by automatically assigning tickets to the employees responsible, for instance. If an ITOA system learns, for example, that more performance is needed at certain times, cloud resources can be activated automatically.

Furthermore, system downtimes can be forestalled. If, for example, the ITOA system notices that the temperature in the data center is rising and also that the physical server's CPU utilization is high (thus generating a lot of heat), it may be necessary to upgrade the air conditioning. Such an action can be initiated in this step. Another example is storage media that are filled with data from the servers until no more memory is available. This too can cause a system to break down. The ITOA system can predict when a hard disk will likely be full and can order a new one in good time, so that a system failure can be prevented.

The following should be completed in this phase:

- Recommendations for action based on analyses
- Recommendations about implementing the next maturity level and optimizing the ITOA architecture
- Training the users to work with the ITOA system, if necessary

Conclusion

This white paper shows how, using IT operations analytics, expensive downtime can be avoided, troubleshooting can be accelerated, and the overall performance of IT can be improved. An ITOA system offers a single point of truth and thus enables searching throughout all IT operating data. Thanks to modern analysis and forecasting methods such as machine learning, failures are predicted and proactive measures can be taken to avoid them.

An ITOA architecture usually comprises one software agent that collects the log, performance, or application data, a log collector or database where this data is stored in a structured way, and a log server that enables the analysis of the distributed data. Different databases and architectures are used depending on the provider.

There are currently around 90 ITOA providers and products on the market. HP Operations Analytics, IBM Operations Analytics, and Splunk have proven to be good general ITOA tools. To find out which tool best meets your needs, we recommend performing a detailed evaluation. The most important criteria here are ease of use, costs, the data sources supported, the adaptability of dashboards and reports, and the quality of visualizations.

What's more, ITOA projects are highly complex, so they should only be performed using a defined procedure model. The six phases – definition of the strategy and goals, as-is analysis and design, implementation and connection of the data sources, data analysis, modeling and evaluation, and optimization and transformation – have proven their worth in practice. Working according to this model offers a structured guideline and means that even complex big data analytics projects can be completed successfully.

Contact us //

mayato GmbH
Am Borsigturm 9
13507 Berlin

info@mayato.com

+49 30 4174 4270 0