Wolfie Christl

EMPLOYEES AS RISKS

A case study on intrusive surveillance and behavioral profiling for cybersecurity, insider risk detection and “compliance”

GEFÖRDERT DURCH Digifonds AK WIEN

A CASE STUDY BY CRACKED LABS

Vienna, August 2024

This publication is part of the project “Surveillance and Digital Control at Work”: crackedlabs.org/data-at-work

Employees as Risks

A case study on intrusive surveillance and behavioral profiling for cybersecurity, insider risk detection and "compliance". Cracked Labs, August 2024.

Author: Wolfie Christl
Edited by: Mike Holohan

Cracked Labs – Institute for Critical Digital Culture
Gumpendorfer Straße 63b, 1060 Vienna, Austria
https://crackedlabs.org

This publication is part of the project "Surveillance and Digital Control at Work", which aims to explore and examine how companies use personal data on workers in Europe.
https://crackedlabs.org/data-at-work

The production of this publication was supported by the "Digitalisierungsfonds 4.0" of AK Wien.

Contents

Summary

Organizations use increasingly intrusive digital monitoring and behavioral profiling to prevent cyberthreats, data leaks and other information security incidents. Employees are seen as major risks. They may enable cyberattacks through carelessness or negligence, for example, by falling victim to a phishing attack or sending information to their private email address, or they may become "insider threats" who intentionally plan to harm the employer. In recent years, organizations have begun using software that analyzes large amounts of activity log records and communications data for purposes that go well beyond cybersecurity. A variety of software systems promise to help them prevent employee misconduct, whether it be criminal, negligent, inappropriate or otherwise undesirable. The boundaries between information security, the protection of corporate information, fraud and theft prevention and the enforcement of compliance with regulatory requirements and organizational policies are becoming blurred.

This case study explores, examines and documents how employers can use software that analyzes extensive personal data on employee behavior and communication for cybersecurity, insider threat detection and compliance purposes. To illustrate wider practices, it investigates software for "security information and event management" (SIEM), "user and entity behavior analytics" (UEBA), insider risk management and communication monitoring from two major vendors. First, it looks into cybersecurity and risk profiling systems offered by **Forcepoint**, a software vendor that was until recently owned by the US defense giant Raytheon. Second, it investigates in detail how employers can use cybersecurity and risk profiling software sold by **Microsoft**, whose "Sentinel" and "Purview" systems provide SIEM, UEBA, insider risk management and communication monitoring functionality. Combined, these systems can **monitor everything employees do or say**, profile their behavior and single them out for further investigation. Similar to predictive policing technologies, they promise not only to *detect* incidents but to *prevent them before they occur*. While organizations can use these software systems for legitimate purposes, this study focuses on their potential implications for employees.

Note: The research in this case study refers to products offered by Forcepoint up until late 2023 (see section 3).

Based on a detailed analysis of software documentation and other corporate sources, this case study documents a wide range of data practices. Both Forcepoint and Microsoft provide far-reaching surveillance capabilities:

- **Monitoring employee behavior and communication.** The systems examined in this study can monitor how employees access and modify files, how they copy them to the clipboard, the applications they use, the websites they visit and their searches, their email and chat conversations, voice calls, video meetings, how they physically access buildings and offices, their performance reviews and even keyboard and screen activity.
- **Analyzing extensive personal data across the organization.** Data from employee computers is accessed via anti-virus, device management or extra monitoring software installed on their devices. Activity logs from almost any enterprise software system used in an organization provide additional information on employee behavior, from Microsoft 365 and Teams to Zoom, Salesforce, Oracle and SAP. Data sources can also include networks, firewalls, spam and web filtering software, badging systems and HR software such as Workday.
- **Singling out suspicious employees and ranking them by risk.** Both Forcepoint and Microsoft offer to continuously calculate risk scores for employees, assess their behavior, rank them by risk and raise alerts about those who are considered potential "insider threats" or otherwise suspicious.
- **Detecting "anomalous" behavior.** Several systems examined in this study promise to "learn" over time how employees usually behave and then try to identify "anomalous" behavior. This AI-based profiling relies on the ongoing analysis of data on past activities of employees across departments and entire organizations.

- **Intrusive inferences and assessments.** Based on data on behaviors and communication, Forcepoint offers to assess whether employees are in financial distress, show "decreased productivity" or plan to leave the job, how they communicate with colleagues and whether they access "obscene" content or exhibit "negative sentiment" in their conversations. Microsoft promises to detect "insider threats" based on assessments about "risky browser usage" and "offensive language". It suggests focusing on employees with a "predisposition" to "violate company policies" and specifically targets "disgruntled employees" who received "poor performance reviews", were demoted, put on "performance improvement plans" or are to be terminated. Organizations can detect almost any type of behavior based on custom data sources, risk indicators and AI-based policies.
- **Pervasive communication surveillance.** Microsoft offers to scan email and chat conversations, voice calls, meeting transcripts and file contents for a wide variety of purposes ranging from "acceptable use" to compliance, cybersecurity and criminal misconduct. Its communication monitoring system promises to detect "profanity", "offensive language", "inappropriate text", threats, harassment and discrimination but also corporate sabotage, data leaks, bribery, money laundering, insider trading, conflicts of interest and "workplace collusion". Employers can receive alerts when certain keywords are mentioned. They can "train" custom AI-based classifiers by providing text samples that represent the type of content they want to detect. Via third-party software, the system can access data from mobile devices, including calls and encrypted messages in WhatsApp or Signal.
- **Investigating past employee activities and screen recordings.** Organizations can use the insider risk and communication monitoring systems examined in this study to further investigate suspicious employees and their past behavior, including their website visits, file and application usage, badging activity and communication contents. For "forensic" investigations, employers can access screen recordings and fine-grained user interaction data on typing activity, clipboard usage or the currently active window at a certain point in time. Forcepoint promises to provide an "over-the-shoulder view" of the employee's computer.
- **Combining cybersecurity and risk surveillance.** As Microsoft's cybersecurity software "Sentinel" can process alerts about suspicious employees from all the other risk profiling systems, it can become a combined security and risk surveillance system. It can analyze millions of log records per second and access up to seven years of past data. Sentinel offers to detect "non-routine actions" and "non-compliant practices" including "insider threats". It promises to help organizations understand whether a suspicious user is a "disgruntled employee who just got passed over for a promotion". Organizations can put certain employees on "watchlists" and perform "dragnet" searches for certain behaviors according to various criteria in real time.
- The study briefly examines other Microsoft technologies for auditing, "data loss prevention" (DLP) and "eDiscovery" and systems from two other vendors. **IBM** offers SIEM, UEBA and insider risk systems similar to Forcepoint and Microsoft. Its communication monitoring system promises to assess "emotions". **Teramind** provides intrusive surveillance software that openly combines security, risk and productivity monitoring.

Organizations must protect themselves from cyberattacks, data loss and criminal misconduct. This is not optional, and, in several ways, mandated by law. Nevertheless, **intrusive security and risk surveillance raises serious concerns** about misuse by employers, disproportionate monitoring and profiling across purposes, flawed risk assessments and arbitrary suspicions. As discussed in the final section of this study, **employers can potentially misuse** these technologies to spy on employees, target organized labor, suppress internal dissent, apply excessive behavioral policing or impose arbitrary disciplinary action. These systems **put employees under general suspicion** and can undermine privacy, human dignity, autonomy, freedom of expression and trust in the workplace. When employees

with “poor” performance reviews receive extra scrutiny, employers can apply more rigid performance monitoring. Surveillance generally increases the power and information asymmetry between organizations and employees.

Employers can widely **customize** the systems provided by Forcepoint and Microsoft. They can either limit or expand data sources and profiling capabilities and apply them either to only a few employees with access to sensitive resources or to their entire staff. They can implement more or less effective **safeguards** such as pseudonymization, access control and auditing. While employers are primarily responsible for deploying these systems, **software vendors influence and shape how they are used**. Forcepoint, whose behavioral surveillance technology was initially funded by the CIA, recommends that organizations implement intrusive profiling across all employees and suggests targeting “internal activists”. Its customers include businesses in all sectors, including in Europe. Microsoft provides similar technology, which is easily available to many employers who already use Microsoft software. As this investigation shows, Microsoft recommends that customers monitor all employee communication at least for “harassment or discrimination detection” and systematically incentivizes them to expand risk surveillance. The findings of this case study suggest that the cybersecurity and risk profiling systems offered by Forcepoint, Microsoft and other vendors help **normalize pervasive employee surveillance** and contribute to its expansion.

The findings will be incorporated in the main report of the ongoing project “Surveillance and Digital Control at Work” (2023-2024), led by Cracked Labs, which explores how companies use personal data on workers in Europe.

1. Introduction, overview and scope

Organizations increasingly use software systems that utilize intrusive surveillance and behavioral profiling to address a wide range of challenges from cybersecurity to "insider threats" to compliance with corporate policies. This case study investigates how these systems process personal data on employees and how organizations can use them, with a focus on the potential implications for employees.

To illustrate wider practices, it examines software for SIEM, UEBA, insider risk management and eDiscovery offered by Forcepoint and Microsoft. While Forcepoint is a pure cybersecurity vendor that is affiliated with the defense and intelligence sector, the cybersecurity and risk profiling systems provided by Microsoft are easily available for both large and smaller organizations who already use Microsoft software. In addition, this case study briefly examines software provided by other vendors. Building on previous German-language research (Christl, 2021) and a literature review, it aims to identify, examine and document data practices that affect workers, based on an analysis of publicly available corporate sources such as software documentation, training videos and marketing materials.

- **Section 2** gives an overview of relevant technologies and software vendors that offer systems for "security information and event management" (SIEM), "user and entity behavior analytics" (UEBA), "insider risk management" (IRM), "data loss prevention" (DLP), communication monitoring and eDiscovery. It addresses how these systems consider employees as risks that require behavioral policing and how organizations can potentially misuse these technologies to spy on employees, suppress internal dissent or target organized labor. It briefly explores how a number of software vendors are affiliated with the defense and intelligence sector.
- **Section 3** investigates cybersecurity and risk profiling software sold by Forcepoint, a major US cybersecurity vendor which was until recently owned by the defense giant Raytheon. Forcepoint's UEBA and insider risk systems monitor employee behavior and communication based on extensive data from devices and log records. They provide intrusive profiling functionality for very different purposes and promise to detect "anomalous" and otherwise suspicious activities. Organizations can further investigate employees and their past behavior including web, application, file, email, chat, call, badging, keyboard and screen activity.
- **Section 4** explores similar technology provided by the enterprise software giant Microsoft. It focuses on Microsoft's SIEM, UEBA, insider risk and communication monitoring systems that are part of two higher-level software systems which Microsoft refers to as "Sentinel" and "Purview". Deeply integrated with Microsoft 365, these systems can monitor employee behavior and communication based on data from many sources, including from devices and from other enterprise software systems such as Zoom, Salesforce, Oracle and SAP. Similar to Forcepoint, Microsoft provides intrusive profiling functionality for very different purposes and promises to detect "anomalous" and otherwise suspicious activities. This section investigates in detail how Microsoft's security and risk profiling systems process personal data on employees and how they can share information with each other. It briefly examines Microsoft's "audit log" functionality and its DLP and eDiscovery systems, the latter of which allows organizations to search for employee information and compile dossiers on them. In addition, it discusses the available safeguards that promise to address employee privacy, data protection and misuse.
- **Section 5** briefly investigates systems from other vendors. This includes SIEM, UEBA, insider risk management and communication monitoring systems provided by IBM and an extremely intrusive employee surveillance system provided by Teramind, which openly combines risk and productivity monitoring.

- **Section 6** summarizes the findings. It presents an overview of data practices identified and documented in this investigation and outlines how the systems provided by Forcepoint and Microsoft process personal data on employees, including processing activities, data sources and purposes.
- **Section 7** reflects on the findings and discusses potential implications for employees.

This case study is part of a series of case studies on systems that process data at the workplace, which are, in turn, part of the **ongoing project**, "Surveillance and Digital Control at Work",[1] led by Cracked Labs. The project aims to explore how companies use personal data on and against workers in Europe, together with AlgorithmWatch, Jeremias Prassl (Oxford), UNI Europa and GPA, funded by the Austrian Arbeiterkammer. The case studies build on **previous research** on the topic (Christl, 2021). They aim to document technologies and data practices by reviewing existing literature and by **examining technologies and software systems** that are available on the market based on publicly accessible vendor information. This includes software documentation and marketing materials, which might be ambiguous and incomplete. Every effort has been made to accurately interpret these corporate sources, but we cannot accept any liability in the case of eventual errors. Where the case studies rely on the examination of corporate sources, it remains largely unclear how employers actually implement, customize and use the functionality provided by these systems. The products and services offered by the software vendors examined in this case study may have changed since the time of the examination. The findings of the case studies will be incorporated into the **main report** of the ongoing project, which will draw further conclusions from the findings.

[1] https://crackedlabs.org/en/data-at-work

2. Enterprise software systems for cybersecurity, compliance and the detection of "insider threats"

Today's corporate IT infrastructure is exposed to numerous cybersecurity threats. The potential damages can get entire organizations into trouble and indirectly affect other groups, for example, when data falls into the wrong hands or critical infrastructure fails. In order to counter these threats, detect cyberattacks and protect an organization's assets, cybersecurity systems increasingly process large amounts of log data and other records from different enterprise software systems used across the organization (see e.g. Vacca, 2017; Riesenecker-Caba, 2023). As this case study shows, this often includes extensive personal data about employees, their behaviors and communication.

2.1 Software for SIEM, UEBA, insider risk management and eDiscovery

Enterprise information security is a broad field. Organizations use a wide range of software that aims to protect networks, devices, applications, databases and information (Vacca, 2017). In addition, organizations have started to use software systems that promise to protect them from diverse risks resulting from criminal, unlawful, negligent, non-compliant, careless, inappropriate or otherwise undesirable employee behavior (Kuldova and Nordrik, 2023; Gelles, 2016). This case study focuses on cybersecurity and risk profiling systems that process extensive data on employee behavior and communication, aim to detect activities that are considered suspicious and provide functionality to further investigate employees. Software vendors use a variety of terms and abbreviations to refer to different types of cybersecurity and risk profiling systems that provide sometimes overlapping functionality.

Systems for **security information and event management (SIEM)** aggregate and combine large amounts of activity log data from multiple and diverse sources including data from an organization's network infrastructure, devices, applications, cloud services and other cybersecurity systems such as anti-virus software. SIEM systems can process millions of activity records in real time and often provide additional functionality for analysis and threat detection (Vacca, 2017; González-Granadillo and González-Zarzosa, 2021). Systems for **user and entity behavior analytics (UEBA)** explicitly analyze employee behavior to detect activities that are considered a potential threat to an organization based on a mix of rules and AI technology. They "learn" over time how users and employees typically behave, try to detect unusual and "anomalous" behavior, constantly calculate risk scores for users and single out suspicious employees (Khaliq et al., 2020; Cardoso, 2021). As such, they aim to prevent cyberattacks and other threats before they occur. UEBA systems monitor a wide range of data sources, systems and employee activities. As this case study shows, this can include data from employee devices (e.g. login, file, application and web browsing activity) and from an organization's network infrastructure and information on employee communication, badging and even performance reviews from the HR system.[2] Many cybersecurity systems provide both SIEM and UEBA functionality. Software vendors include, for example, Microsoft, IBM, Splunk (Cisco), Exabeam, Securonix, Forcepoint, Gurucul, LogRhythm, Rapid7, Fortinet, CrowdStrike and Micro Focus (Gartner, 2018; Gartner, 2022).

While SIEM and UEBA systems address both internal and external threats, systems for **insider risk management (IRM)** explicitly focus on the detection of "insider threats", i.e. employees who are considered a threat to the organization. They use behavioral monitoring and profiling to single out suspicious employees based on similar records on employee activity as SIEM and UEBA systems. In addition to file, application and web browsing activity,

[2] See section 6

this specifically includes printing and USB activity. They typically provide functionality for intrusive communication and device surveillance, from monitoring employee conversations via email, chat, voice call and video meeting to recording screen activity or even keystrokes (Gartner, 2020). As this case study shows, Microsoft's insider risk system specifically targets "disgruntled" employees who received "poor performance reviews", were demoted or put on "performance improvement plans", conducted previous "policy violations" or are to be terminated, based on data from HR systems. It can also utilize data on visits to "inappropriate" websites or the use of "offensive" language as indicators of potential insider risks.[3] Some systems make intrusive assessments based on behavioral profiling by trying to detect, for example, whether employees show "decreased productivity", are in "financial distress" or plan to leave the job.[4] They may also include external data about employees, from criminal records and "background check" information[5] to data on social media activity (Gelles, 2006). Insider risk systems often include UEBA functionality or they are part of SIEM or UEBA systems. Software vendors include Microsoft, Forcepoint, DTEX, Exabeam, Splunk (Cisco), Gurucul, LogRhythm, Rapid7, Fortinet, Micro Focus, Code42, Proofpoint, Veratio, Teramind and ActivTrak (Gartner, 2020). Several systems also include **data loss prevention (DLP)** functionality, which aims to prevent customer data, trade secrets and other sensitive information from leaving an organization's IT infrastructure (Alneyadi et al., 2016).

Other systems focus on **communication monitoring and profiling** for both security and compliance purposes. While there is no established terminology for these systems, they typically scan employee conversations via email or chat to detect inappropriate communication activities in order to meet regulatory requirements, enforce internal policies or detect security threats. This typically includes monitoring file content and, based on automated transcripts, even conversations in voice calls and video meetings. After suspicious content is detected based on keywords or AI-based classifiers, the corresponding employees can be further investigated. The market research firm Gartner, which refers to systems for "digital communications governance", lists several software vendors such as Microsoft, Proofpoint, Barracuda, Symphony, TeleMessage, CellTrust and Veritas (Gartner, 2023). Microsoft's "communication compliance" system can detect very different types of communication content including "inappropriate" or "offensive" language, discrimination, harassment, threats, bribery, gift exchanges, money laundering, insider trading, "workplace collusion", corporate sabotage and data leaks.[6] It can be integrated with Microsoft's SIEM and insider risk systems.[7] Some systems can record and monitor voice calls, SMS and encrypted messages (e.g. WhatsApp, Signal) directly from the employees' mobile devices. TeleMessage offers to collect data on phone calls and SMS via partnerships with mobile carrier networks.[8]

While communication monitoring systems analyze employee conversations and document contents in real time or almost real time, systems for **eDiscovery** offer to identify, analyze and preserve all communication contents, documents and other "electronic information" stored across an organization that is relevant to certain internal or external investigations, including against employees. They typically help to constantly collect and archive relevant information and allow broad searches that potentially affect a large number of employees and other persons (Sachowski,

[3] See section 4.1
[4] See section 3
[5] Ibid.
[6] See section 4.2
[7] See sections 4.3 and 4.8.1
[8] See section 4.7

2018; Gartner, 2015). Vendors that provide eDiscovery software include Microsoft, Relativity (formerly kCura), OpenText (formerly Recommind), ZyLab and Exterro (Gartner, 2015).

2.2 Spying on "disgruntled employees", "internal activists" and "organized labor"?

Employees are increasingly seen as potential cybersecurity risks and put under general suspicion. As shown in Figure 1 (left), the technology giant Intel suggests in a "white paper" about enterprise security that employers should consider workers as "threats" in many respects. Employees may act or work in the interest of hostile competitors, suppliers, partners, nation states, organized crime or even terrorism. They may be simply careless, distracted or poorly trained or they may aim to intentionally harm the company – whether they are motivated by profit, because they are dissatisfied at work or because they are "activists", i.e. "highly motivated supporters of a cause".

THREAT AGENT TYPE	DEFINITION
Reckless Insider	Person who knowingly and deliberately circumvents safeguards for expediency but does not intend harm or serious consequences
Untrained/Distracted Insider	Person with harmless intent who inadvertently misuses systems or safeguards
Outward Sympathizer	Person who knowingly misuses the enterprise's systems to attack others in support of a cause external to the enterprise, but with harmless intent to the enterprise itself
Supplier	Business partner who seeks inside information for business advantage over its own competitors (that is, other suppliers)
Partner	Business partner with whom the enterprise has voluntarily shared sensitive data for collaborative efforts and who may either accidentally or deliberately expose that information
Irrational Individual	Person acting with illogical purpose and behavior
Thief	Opportunistic person with profit motive
Disgruntled Insider	Unhappy current or former insider with intent to harm the enterprise, industry, or fellow insider
Activist	Highly motivated supporter of a cause who does not engage in physical violence
Terrorist	Person who relies on physical violence or extreme acts to support a socio-political agenda
Organized Crime	Crime syndicate with significant resources and attack skills
Competitor	Business adversary who competes for customers, revenues, public exposure, or resources
Nation State	State-sponsored attacker with significant resources, and able to affect a major disruption to even national scale

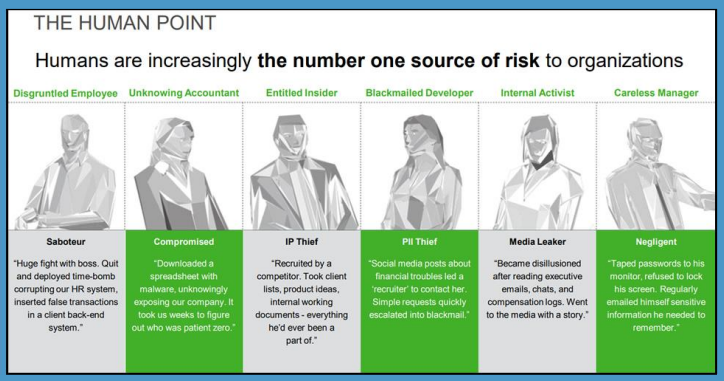

Figure 1: Employees as "insider threats" (Intel, Forcepoint)[9]

The cybersecurity vendor Forcepoint refers to "humans" as "the number one source of risk to organizations". As Figure 1 (right) from a corporate presentation shows, possible threats include employees with compromised access credentials who downloaded malware, negligence and data theft but also "disgruntled employees" who had a "huge fight with the boss" and "internal activists" who leak information to journalists. It is not difficult to imagine employers misusing the extensive surveillance capabilities provided by today's cybersecurity and risk monitoring systems to inappropriately or illegally spy on employees or even on worker representatives and works councils.

A job advertisement published by Amazon on its website in 2020 demonstrates a **worst-case scenario**. As Figure 2 shows, Amazon was hiring an "intelligence analyst" for the "Global Intelligence Program" of its "Global Security Operations" unit who would work on "sensitive topics that are highly confidential, including labor organizing threats against the company" and help spy on "organized labor, activist groups [and] hostile political leaders". Required qualifications included "experience working with global risk intelligence, incident response, large data analytics software", preferably in the "intelligence community, the military, law enforcement, or a related global security role in the private sector". The analyst was expected to "learn and understand a broad range of Amazon data resources",

[9] Figures © Intel, Forcepoint. The figures serve as basis for the discussion of the corporate practices examined in this study. Sources: Intel (2015), A Field Guide to Insider Threat. IT@Intel White Paper, p. 6: https://www.researchgate.net/publication/324068663_A_Field_Guide_to_Insider_Threat_Understanding_Insider_Threat_Vectors_to_Further_Improve_Enterprise_Security_Strategies; Forcepoint (2017), Forcepoint UEBA. User & Entity Behavior Analytics, p. 15: https://www.ctc-g.com.sg/wp-content/uploads/2021/12/Forcepoint-UEBA-User-Entity-Behavior-Analytics.pdf [30.10.2023]

"exploit these datasets for new insights" and "monitor various collection platforms for incidents that pose direct and indirect risk to Amazon".

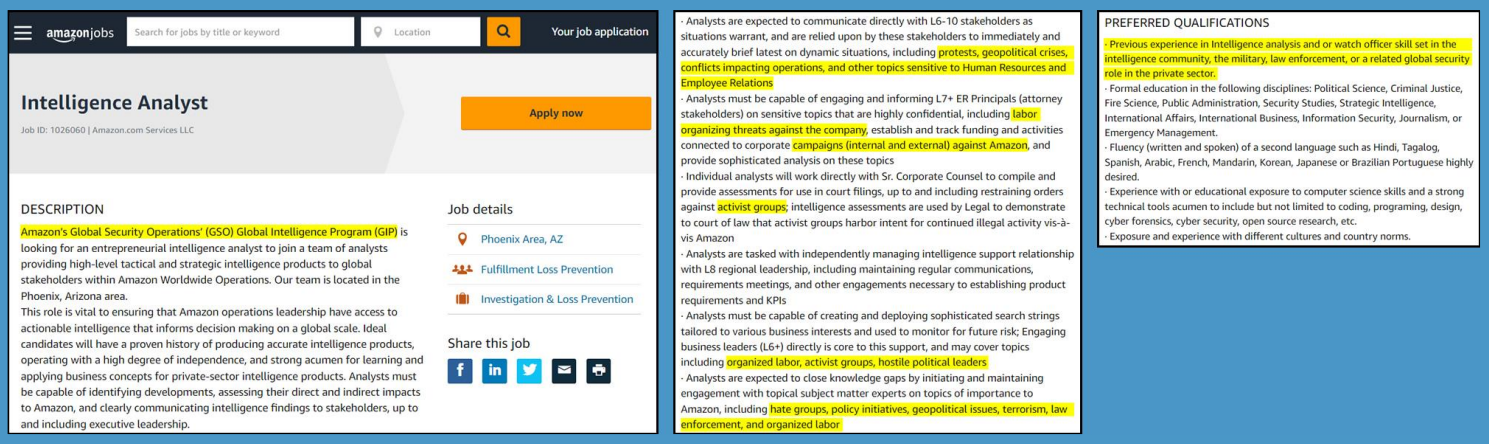

Figure 2: Hiring an "intelligence analyst" who would spy on "organized labor" and "activist groups" (Amazon)[10]

After the author of this study pointed to this job advertisement on Twitter, Amazon silently took it offline and published another version of it, with the most problematic phrases removed.[11]

2.3 Security and risk profiling across purposes as "predictive policing"

Even if organizations do not openly misuse cybersecurity and risk profiling systems to illegally spy on employees, the scale and depth of data collection raises serious concerns about the privacy-intrusive and disproportionate use of extensive personal data about employees for different purposes.

Joseph Da Silva (2022) argues that the perception of cybersecurity threats as permanent, existential and potentially catastrophic can **justify and normalize intrusive surveillance** that would otherwise not be accepted and which could be used beyond the purposes of cybersecurity. He conducted semi-structured interviews with executives who are responsible for information security in organizations, who describe difficulties in achieving "a balance between inspection and surveillance" and see themselves as performing a "policing" function. Consequently, he found that the cybersecurity departments in these organizations "appeared to function as an official police force".

By expanding the scope of organizational risk surveillance from cyberattacks to "insider threats" to non-compliance with internal and external policies, "everyone—from employees, customers, clients, to third parties—is by default deemed a potential threat, risk, and criminal", as Tereza Østbø Kuldova (2022) writes in her book on the "compliance-industrial complex". Employees and other actors are "seen primarily as suspects and as guilty until proven innocent, where this innocence can at any moment flip into the opposite and therefore requires permanent monitoring and surveillance". She presents Microsoft's risk and compliance management system Purview, which is further examined in section 4 of this case study, as an example of the "progressive integration of the risk and compliance universe into single governance platforms" that monitor employees for very different purposes, including data leaks,

[10] Figures © Amazon. The figures serve as basis for the discussion of the corporate practices examined in this study. Source: Amazon Jobs website, 1.9.2020, archived at: https://web.archive.org/web/20200901125228/https://www.amazon.jobs/en/jobs/1026060/intelligence-analyst [30.10.2023]
[11] https://twitter.com/WolfieChristl/status/1300776980602925060

IP theft, fraud, conflicts of interest, policy violations, harassment and "inappropriate" language. Because these systems move towards the pre-emptive detection of "anomalies" and behaviors that are seen as indicators for future breaches, they "increasingly imitate predictive policing", according to Kuldova, which is why she considers them to be part of an "operating system of a pre-crime society". She also points to the potential ramifications of algorithmic monitoring, from flawed inferences to inaccurate risk assessments, and suggests that going after employee fraud at the individual level typically relies on "denial of systemic fraud". Consequently, she generally questions the idea "that it is individual 'bad apples' that are responsible for fraud and that it is the 'rogue' employee that needs to be identified and surgically removed in order to not spoil the corporate body".

The findings of this case study suggest that the boundaries between cybersecurity, fraud and theft prevention, the protection of confidential information and the enforcement of "compliance" with laws, guidelines, policies, codes of conduct and other organizational rules are becoming increasingly blurred.

2.4 Cybersecurity software and the defense and intelligence sector

As cybersecurity has long been considered a "cornerstone" of national security and defense (see e.g. Kovacs, 2018), it is not surprising that a number of enterprise cybersecurity software vendors are closely affiliated with the defense and intelligence sector. This is specifically true for systems that are examined in this case study, which use behavioral profiling based on large amounts of data.

The **CIA's venture capital firm** In-Q-Tel[12] has been and continues to be an investor in several cybersecurity companies,[13] for example, in the UEBA vendor Interset,[14] which was acquired by Micro Focus in 2019,[15] and in RedOwl,[16] which was acquired by Forcepoint in 2017 and is now known as Forcepoint UEBA.[17] The major cybersecurity vendor Forcepoint itself was, until 2020, owned by the US defense giant Raytheon, before it was sold to a private equity firm.[18] A co-founder of the cybersecurity firm RedOwl which has now become Forcepoint's UEBA system is a former US army intelligence and NSA officer who was previously the CEO of Berico Technologies, which was involved in a large-scale plan to **discredit labor unions** in the US that was uncovered in 2011 and which included data gathering from social media.[19]

Cybersecurity software often includes far-reaching surveillance technology, and while the affiliations between some vendors and the security state are not surprising, organizations in many industries exposing their employees to military-grade surveillance raises concerns.

[12] https://www.cia.gov/stories/story/cia-contributions-to-modern-technology-75-years/ [2.5.2024]
[13] https://www.iqt.org/portfolio/ [2.5.2024]
[14] Ibid.
[15] https://www.microfocus.com/about/press-room/article/2019/micro-focus-completes-acquisition-of-interset-to-further-expand-cyber-security-expertise/ [2.5.2024]
[16] https://www.iqt.org/portfolio/ [2.5.2024]
[17] https://web.archive.org/web/20220305113412/https://www.forcepoint.com/landing-page/redowl [2.5.2024]
[18] https://www.wsj.com/articles/raytheon-takes-control-of-forcepoint-cybersecurity-business-11580385744, https://www.crn.com/news/security/raytheon-unloads-security-subsidiary-forcepoint-to-private-equity [2.5.2024]
[19] https://www.motherjones.com/politics/2011/02/chamberleaks-strategies-defame-foes-us-chamber-revealed/, https://archive.thinkprogress.org/chamberleaks-military-contractors-palantir-and-berico-under-scrutiny-480cae35e353/, https://cryptologicfoundation.org/about/governance.html/title/member-mr-guy-filippelli [2.5.2024]

3. Forcepoint's UEBA and "insider threat" systems

Forcepoint, formerly known as Websense,[20] is a major US cybersecurity vendor.[21] Until 2020, it was owned by the US defense giant Raytheon and then sold to a private equity firm.[22] The company's customers include both government agencies and large defense contractors, banks, insurance firms, healthcare providers, manufacturers, oil and gas companies, utility providers, telecommunication firms, airlines, hotel chains and supermarket chains. While many of them are unnamed, Forcepoint names customers such as Boeing, Qualcomm, Toyota, Honda, Walmart, Burger King, Swisscom and Swiss Life.[23] In late 2023, Forcepoint's government cybersecurity business, which primarily serves the US government, federal agencies and government contractors, was sold to another private equity firm for $2.45 billion and rebranded as "Everfox". The company's commercial cybersecurity business continues to operate under the brand name Forcepoint.[24]

Note: The research for Section 3 on Forcepoint's UEBA and insider threat systems was primarily conducted between 2021 and 2023, with additional corporate sources reviewed in early 2024. In late 2023, Forcepoint's government cybersecurity business was sold to a private equity firm and rebranded as "Everfox". As of August 2024, Everfox appears to be the entity selling the behavioral analytics and insider threat systems examined in this case study. Forcepoint continues to offer a broad range of other cybersecurity products. Therefore, **the research in this case study refers to products offered by Forcepoint up until late 2023**, or as indicated by the dates of the corresponding sources in the footnotes.

Forcepoint sells **enterprise cybersecurity software** for devices, networks and cloud environments.[25] Its firewall and email security products offer functionality for filtering and blocking access to websites[26] and for analyzing and filtering email contents.[27] The company also offers software for data loss prevention (DLP),[28] behavior analytics (UEBA)[29] and a system that promises to detect "insider threats".[30] The following sections focus on how Forcepoint's UEBA and insider threat systems process personal data about employees. Organizations can use these systems for extensive behavioral monitoring and profiling.

3.1 Employee profiling and risk scoring based on extensive personal data

Forcepoint's **UEBA system**, which the company also refers to as "Forcepoint Behavioral Analytics", continuously analyzes a wide range of log data about employee activities in order to calculate ongoing risk assessments based on "big data analytics and machine learning".[31] It evaluates information about employees' login activities, the programs they use on their computer, the files they access, modify or move, the websites they visit, their Google searches and any communication via email, chat or phone. This can include the analysis of communication contents, in the case

[20] https://www.zdnet.com/article/raytheon-websense-rebrands-as-forecepoint-acquires-intel-securitys-stonesoft/ [5.2.2024]
[21] See e.g. Forrester (2023): The Forrester Wave: Data Security Platforms, Q1 2023; Gartner (2028): 2018 Magic Quadrant for Secure Web Gateways; Gartner (2023): 2023 Gartner Magic Quadrant for Security Service Edge (SSE); Gartner (2023): 2023 Gartner Magic Quadrant for Single-Vendor SASE; Gartner (2020): 2020 Gartner Magic Quadrant for Enterprise Data Loss Prevention; Gartner (2019): 2019 Magic Quadrant for Network Firewalls
[22] https://www.wsj.com/articles/raytheon-takes-control-of-forcepoint-cybersecurity-business-11580385744, https://www.crn.com/news/security/raytheon-unloads-security-subsidiary-forcepoint-to-private-equity [5.2.2024]
[23] https://www.forcepoint.com/company/customers [5.2.2024]
[24] https://www.reuters.com/markets/deals/tpg-buy-forcepoint-unit-francisco-partners-2023-07-10/, https://www.forcepoint.com/newsroom/2023/tpg-completes-acquisition-forcepoint-global-governments-and-critical-infrastructure, https://www.businesswire.com/news/home/20240130753876/en/Forcepoint-Federal-Rebrands-as-Everfox-to-Reflect-New-Era-of-Defense-Grade-Cybersecurity [5.2.2024]
[25] https://www.forcepoint.com/products [5.2.2024]
[26] https://www.forcepoint.com/product/secure-web-gateway-swg, https://www.forcepoint.com/cyber-edu/web-content-fitlering [5.2.2024]
[27] https://www.forcepoint.com/product/email-data-loss-prevention-dlp, https://www.forcepoint.com/cyber-edu/secure-email-gateway [5.2.2024]
[28] https://www.forcepoint.com/product/dlp-data-loss-prevention [5.2.2024]
[29] https://www.forcepoint.com/product/ueba-user-entity-behavior-analytics [5.2.2024]
[30] https://www.forcepoint.com/product/fit [5.2.2024]
[31] https://www.forcepoint.com/product/ueba-user-entity-behavior-analytics [5.2.2024]

of phone calls by means of automated transcription of speech into text, and even performance review information from HR systems and badge data about physical visits to offices and rooms.[32] [33] Figure 3 (left) shows a part of the user interface with a ranked list of named employees that are assessed as security risks. For each employee, the system displays a risk score that is calculated based on constant digital profiling.

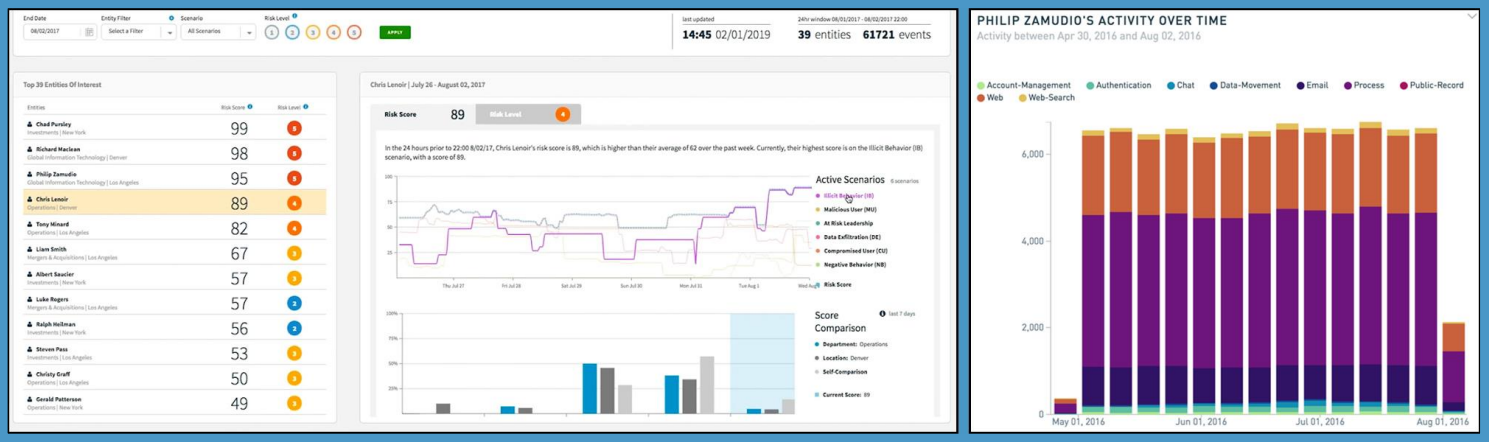

Figure 3: High-risk employees and activity report (Forcepoint)[34]

Figure 3 (right) shows how the system displays a chart that indicates how much time a named employee spent on activities such as using programs, moving files, visiting websites, searching the web and using email over three months.

3.2 Risk models and scores – suspicious employees, "negative" behaviors and "decreased productivity"

The "entities of interest" list, as shown in Figure 3 (left), represents a ranked list of the "most suspicious" employees over the last 24 hours including their names and risk scores.[35] These scores indicate which employees "are behaving different than usual" and thus "may need to be investigated".[36] To calculate the overall risk score for an individual, the system assesses the risk associated with a number of pre-configured risk "scenarios" that are based on **extensive employee profiling** over time.

The chart in Figure 3 (left) illustrates how a particular employee's risk score for the "Illicit Behavior" scenario has significantly increased over the course of a week. In contrast, the score for the "Compromised Account" scenario, i.e. the risk that the employee's account has been hacked, is low. Another chart compares these scores to average values for "peers" in the department and the employee's own behavior in the past. Scenarios are used to "profile activity and anomalies of interest" in order to "detect increasingly risky activity over time".[37] The risk score for a

[32] Forcepoint (2021): Forcepoint Behavioral Analytics User Manual, v3.3.x, 3.4.x, 17.2.2021, p. 1-3, https://www.websense.com/content/support/library/ueba/v33/user_manual/user_manual.pdf [15.2.2024]
[33] https://www.forcepoint.com/sites/default/files/resources/files/brochure_ueba_solution_brief_en.pdf [15.2.2024]
[34] Figures © Forcepoint. The figures serve as basis for the discussion of the corporate practices examined in this study. Sources: video min 1:01: https://www.youtube.com/watch?v=Y3mum4tSmXI, p. 41: https://www.websense.com/content/support/library/ueba/v341/user_guide/FBA_User_Guide_3.4.1.pdf [30.10.2023]
[35] p. 35, https://www.websense.com/content/support/library/ueba/v33/user_manual/user_manual.pdf [15.2.2024]
[36] Ibid., p. 3
[37] Ibid.

scenario is calculated from risk scores for certain behaviors. To detect risky behaviors, Forcepoint uses "behavioral models". The following table shows a selection of pre-built **scenarios and behavioral models**, which Forcepoint recommends setting up for an organization's information security infrastructure:[38]

Scenario	**Behavioral Risk Model**	**Description** (original quotes, unless indicated otherwise)
Data Exfiltration	Internal Data Movement	Anomalous movement of data within the enterprise
	External Data Movement	Anomalous data volumes moving outside of the network
	File Operations	Anomalous interactions with files, such as opening them to check the contents of the file
	File Share Cardinality	Anomalous number of distinct file shares accessed per user
	Data Reconnaissance	Activity conducive with searching for data throughout the enterprise
	Data Loss	Employees who are leaking sensitive information
Suspicious User	Network Reconnaissance	Activity conducive with exploring the network to discover assets of interest
	Suspicious Authentication	Abnormal authentication activity that be conducive with asset discovery
	Access Request	Activity conducive with requests a higher-level of privileges to commit the suspicious action
	Account management	Account management activity by employees who usually do not have any account management activity
	Process Activity	Privileged activity that could facilitate the suspicious actions
	Systems Administration	Activity conducive with […] damaging system configurations
	Code System Components	Interactions with core system files and components
	Suspicious Research	People researching ways to commit suspicious actions
	Physical Access	Abnormal physical access to sensitive areas
Compromised User Account	Phishing	Risk of employees falling victim to phishing, for example through malicious attachments or links in emails [description by the author of this study]
	Malware	Risk that employees have unintentionally installed malware [description by the author of this study]
Negative Workplace Behavior	Obscene Content	Obscene content activity, either through web searches or web browsing
	Negative Sentiment	Negative sentiment and signs of improper discussions within communications
	Flight Risk Communications	Communications conducive with an employee leaving, such as emailing a resume or searching for a new job
	Financial Distress Communications	Communications activity indicative of financial turmoil and employees looking for a way to get resolve the issue
	Oversight Evasion	Signs of oversight evasion attempts within communications
	Decreased Productivity	Employees spending a large amount of time doing non work related tasks
	Corporate Disengagement	Employees who are not interacting with core company assets
	Recipient Cardinality	Users talking to less employees, compared to their previous baselined activity
	Workplace Violence	Communications that could indicate a workplace violence incident
	Sexual Harassment	Communications conducive with a sexual harassment incident
Illicit Workplace Behavior	Corporate Espionage	Employees communicating with competitors while still at their current company, but specifically mentioning their current company [intellectual property]
	Whistleblowing	Activity conducive with a whistleblowing incident, where a person is in contact with media domains and showing signs of willingness to leak company information
	Clearance Evasion	People researching ways to omit security clearance information or ways to deceive a polygraph

Table 1: Employee profiling, risk scenarios and behavioral models (Forcepoint)[39]

[38] Ibid., p. 85ff; source for the "Compromised User Account" scenario: https://www.ctc-g.com.sg/wp-content/uploads/2021/12/Forcepoint-UEBA-User-Entity-Behavior-Analytics.pdf [22.4.2024]
[39] Ibid.

While the detection of phishing, malware, suspicious login activities or unusual changes to system files can be considered part of the field of IT security in a narrower sense, most other behavioral models listed in Table 2 disproportionately interfere with the rights and freedoms of employees and put them under general suspicion by almost completely monitoring their everyday working lives. Forcepoint recommends that monitoring "should include all employees at a company".[40]

Far-reaching behavioral monitoring. The behavioral models associated with the "negative workplace behavior" scenario, as listed in Table 1, aim to identify whether employees are in financial distress, whether they intend to leave the job, how they communicate with colleagues, whether they access "obscene" content or whether there is a "negative sentiment" in their communications. Even data on "decreased productivity", and, as such, data about work performance, is monitored. It is questionable as to whether observing all communication activities is the appropriate means to address workplace violence and sexual harassment. The detection of sophisticated efforts to move data out of an organization's systems may require monitoring activities that involve files and documents. While such a measure may be justified to prevent industrial espionage in some cases, using it to prevent media leaks or other activities that aim to address corporate misconduct is a slippery slope to disproportionate behavioral policing.

3.3 Analyzing and investigating employee activity in detail

While Forcepoint's risk scenarios and models aim to detect behavior that is considered a risk to the organization, employees can subsequently be **placed under special observation** in order to investigate their activities more closely. A promotional video provided by Forcepoint demonstrates how the system is used to analyze the behavior of an airport employee suspected of drug trafficking. Figure 4 shows screenshots from the video.

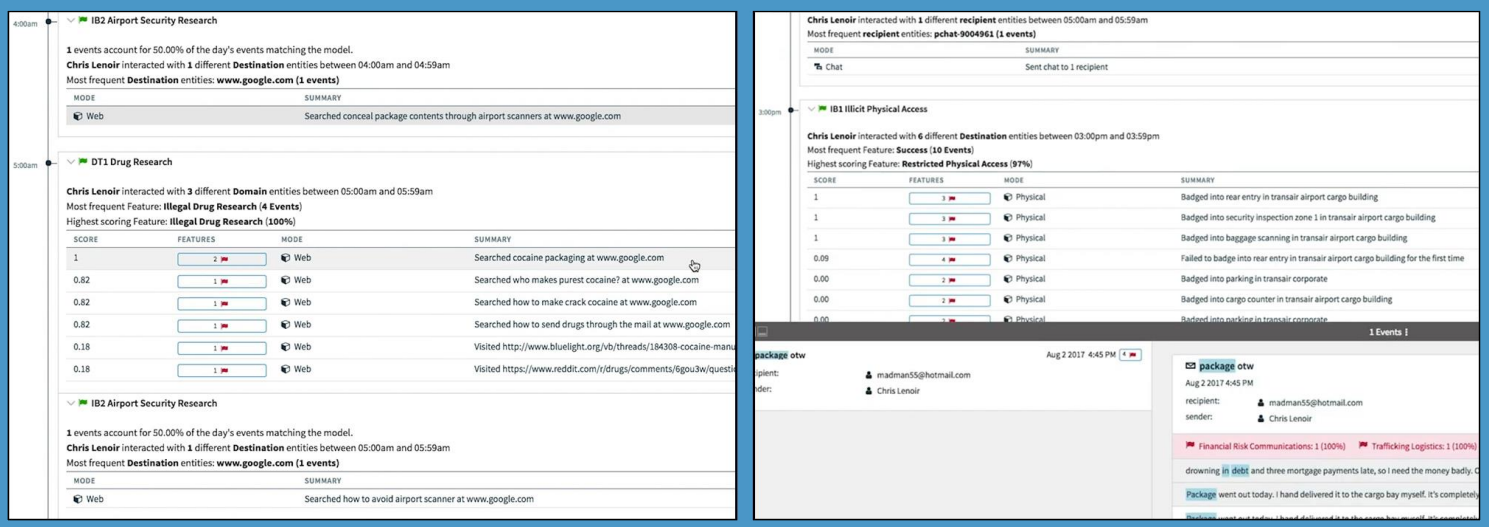

Figure 4: Analyzing activities of an employee who is suspected of criminal misconduct (Forcepoint)[41]

Figure 4 (left) illustrates how the system displays and analyzes a series of website visits and Google searches that are considered suspicious. The employee in question searched for how drugs are processed, packaged and shipped

[40] p. 34, https://www.websense.com/content/support/library/ueba/v33/user_manual/user_manual.pdf [15.2.2024]
[41] Figures © Forcepoint. The figures serve as basis for the discussion of the corporate practices examined in this study. Source: video min 1:27 and 2:02: https://www.youtube.com/watch?v=Y3mum4tSmXI [30.10.2023]

and how they can be smuggled through the airport security check. He also visited two websites on the subject. Later, he visited a number of physical premises on the airport according to badge data and was denied access to one place, as shown in Figure 4 (right). The system also observed him sending a chat message and an email. The content of the email contains words such as "debt" and "package", which triggered the "financial risk communication" and "trafficking logistics" classifications. In this example, the system uses custom risk scenarios for profiling, such as "drug research" and "airport security research", which are tailored to the airport environment.[42] The fact that the observed email involves an external email address demonstrates that the system can monitor both employees and persons outside an organization.

3.4 UEBA data sources, data categories and sensitive inferences

According to Forcepoint, its UEBA system analyzes "data from broad sources" and can "leverage" an organization's "entire IT ecosystem".[43] It can, for example, utilize data from major enterprise software systems provided by Microsoft (Windows, Active Directory, Exchange, Office 365, Skype), Salesforce (CRM, Slack), SAP (Concur), Cisco (networking infrastructure) and Workday (HR). To access data from employee devices, it can utilize "endpoint" monitoring software provided by Forcepoint and other vendors (e.g. Veriato, Digital Guardian). Data sources can also include Forcepoint's data loss prevention (DLP) system, third-party networking and cybersecurity systems (e.g. SIEM, DLP, proxy). In addition to data on communications (email, chat, voice, SMS), devices (file activity, process activity, web activity), authentication (login activity) and networking (proxy and VPN logs), data sources can include print logs, building access logs, "geolocation/GPS logs" and even "performance reviews" from HR systems. Forcepoint also mentions "public records" about "criminal history" or "financial distress".[44] The system can analyze "off hour" and "weekend" activity, which suggests that it potentially monitors the private lives of employees.[45]

Sensitive inferences about communication and web activity. To detect "negative sentiment" in communication activity, employers can maintain customizable lists of keywords, i.e. "terms and phrases that indicate opinions and/or emotions". Forcepoint's "official" keyword list for "negative sentiment" includes, for example, the words "anger", "disappointed" and "mockery". Employers can also maintain customizable lists of websites that, if visited by an employee, may indicate the intent to leave the company.[46]

Forcepoint's "web security" system, which monitors the websites visited by employees and blocks them from accessing sites that contain malware and viruses,[47] groups websites into categories and "risk classes".[48] Website categories that are classified as a potential "legal liability" for employers include, for example, adult material, gambling, file sharing and weapons, as shown in Table 2. Website categories that are classified as an indicator for "productivity loss" include social media and entertainment, but also sites about abortion, health, drugs (including

[42] Ibid.
[43] https://www.forcepoint.com/product/ueba-user-entity-behavior-analytics [22.2.2024]
[44] https://www.forcepoint.com/sites/default/files/resources/files/brochure_ueba_solution_brief_en.pdf, https://www.forcepoint.com/sites/default/files/resources/files/ueba-discover-and-stop-insider-threat.pdf, https://www.forcepoint.com/sites/default/files/resources/datasheets/ueba-platform-architecture-overview_0.pdf [22.2.2024]
[45] p. 19, https://www.websense.com/content/support/library/ueba/v33/user_manual/user_manual.pdf [15.2.2024]
[46] Ibid, p. 51
[47] https://www.forcepoint.com/cyber-edu/web-security-gateway [22.4.2024]
[48] Forcepoint (2022): Forcepoint Web Security, Administrator Help, v8.5.x, last modified April 2022, p. 38ff, https://www.websense.com/content/support/library/web/v85/web_help/web_help.pdf

"abused drugs" and "prescribed medications"), religion and political organizations (including "worker organizations").[49]

Risk class	Selected website categories (original quotes)
Legal Liability	• Adult Material (includes Adult Content, Lingerie and Swimsuit, Nudity, Sex) • Peer-to-Peer File Sharing • Gambling • Hacking, Proxy Avoidance • Intolerance • Militancy and Extremist • Violence • Weapons
Productivity Loss	• Abortion (includes Pro-Choice, Pro-Life) • Sex Education • Entertainment Video, Internet Radio and TV, Peer-to-Peer File Sharing, Streaming Media, Surveillance, Viral Video • Drugs (includes Abused Drugs, Marijuana, Nutrition, Prescribed Medications) • Education (includes Cultural Institutions, Educational Institutions) • Entertainment (includes Media File Download) • Gambling • Games • Political Organizations • Health • Job Search • News and Media (includes Alternative Journals) • Religion (includes Non-Traditional Religions, Traditional Religions) • Shopping (includes Internet Auctions, Real Estate) • Social Organizations (includes Professional and Worker Organizations, Service and Philanthropic Organizations, Social and Affiliation Organizations) • Social Web - Facebook (includes Facebook Apps, Facebook Chat, Facebook Commenting, Facebook Events, Facebook Friends, Facebook Games, Facebook Groups, Facebook Mail, Facebook Photo Upload, Facebook Posting, Facebook Questions, Facebook Video Upload) • Social Web - YouTube (includes YouTube Commenting, YouTube Sharing, YouTube Video Upload) • Society and Lifestyles (includes Alcohol and Tobacco, Blogs and Personal Sites, Gay or Lesbian or Bisexual Interest, Hobbies, Personals and Dating, Restaurants and Dining, Social Networking)

Table 2: Website categories and "risk classes" (Forcepoint)[50]

Processing personal data about employees who visited websites categorized with the label "abortion" and other highly sensitive categories represents intrusive digital profiling.

3.5 Investigating "insider threats", keyboard and screen activity

While the prevention of "insider threats" is only one of several purposes for the UEBA functionality described in the previous sections, Forcepoint also offers software that explicitly addresses "insiders",[51] i.e. employees that are considered risks. The company's **insider threat system** collects even more extensive behavioral data about employees including from Windows and Apple "endpoints",[52] i.e. from their computers and other devices. Forcepoint refers to the system's functionality as "user activity monitoring", which "provides analysts and investigators with deep visibility into all user endpoint activity". It can collect "behavioral data from multiple endpoint channels for

[49] Ibid.
[50] Ibid.
[51] https://www.forcepoint.com/product/fit [26.2.2024]
[52] https://www.forcepoint.com/sites/default/files/resources/datasheets/datasheet-fit-insider-risk-solutions-en_0.pdf, https://www.forcepoint.com/sites/default/files/resources/files/brochure_triton_ap_endpoint_en_0.pdf [26.2.2024]

full context of user activity".[53] As Figure 5 (left) indicates, it can monitor device activity such as the programs used, websites visited, emails sent and received, file operations and even keyboard and clipboard activity.

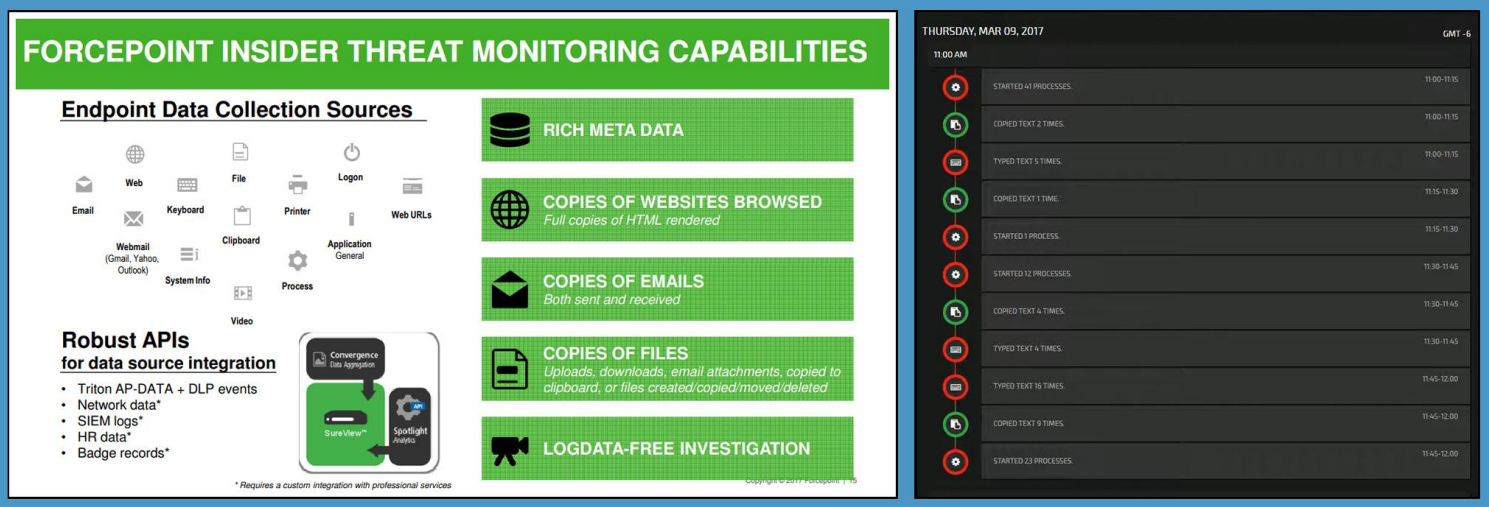

Figure 5: Data sources and employee activity monitoring to detect "insider threats" (Forcepoint)[54]

Forcepoint's insider threat system promises to monitor both "what employees are doing" and "who the employees are". This can include HR information about "performance reviews", "promotions and compensations" and terminations.[55] Similar to Forcepoint's UEBA system, it constantly assesses behaviors and calculates risk scores for employees. The insider threat system enables employers to "automatically identify" the "riskiest users" by "automatically scoring" them, and it provides "forensic evidence" for "investigations, prosecution, and compliance".[56]

Keyboard and screen activity. A product demonstration video shows how the system displays information about "high risk employees" and their past activities.[57] As shown in Figure 5 (right), the user interface can display detailed records about activities performed by a specific named employee on a specific day, including records about typing and copying text on the keyboard. The system can record all user activity from an employee's Windows or Mac OS computer.[58] According to Forcepoint, its "video capture and replay" functionality provides "complete, near-real-time context with an 'over-the-shoulder' view of the end-user's workstation"[59] and "unparalleled visibility into suspicious behaviors".[60] Investigators can "easily call the desktop video replay for high-risk users", allowing for "attribution as well as showing employee intent", which is "admissible in a court of law".[61]

[53] Ibid.

[54] Figures © Forcepoint. The figures serve as basis for the discussion of the corporate practices examined in this study. Sources: presentation, 2017, not available online anymore: https://oldsite.amcham.gr/wp-content/uploads/2017/11/NICK-NICOLESCU.pdf [accessed 18.3.2021], video min 22:05: https://www.youtube.com/watch?v=qSHexqYW-jE [30.10.2023]

[55] https://www.forcepoint.com/sites/default/files/resources/datasheets/datasheet-fit-insider-risk-solutions-en_0.pdf [26.2.2024]

[56] Forcepoint (2016): SureView Insider Threat, https://www.forcepoint.com/sites/default/files/resources/files/brochure_sureview_insider_threat_en.pdf [26.2.2024]

[57] Video min 21:33: https://www.youtube.com/watch?v=qSHexqYW-jE [30.10.2023]

[58] https://www.forcepoint.com/sites/default/files/resources/files/brochure_sureview_insider_threat_en.pdf [26.2.2024]

[59] Forcepoint (2016): SureView Insider Threat, Datasheet, https://www.exit123c.com/wp-content/uploads/2016/05/Forcepoint_sureview_insider_threat_datasheet.pdf [26.2.2024]

[60] https://www.forcepoint.com/sites/default/files/resources/files/brochure_sureview_insider_threat_en.pdf [26.2.2024]

[61] Ibid.

3.6 Normalizing pervasive employee surveillance?

The functions listed in Table 1 for detecting attempts to evade security clearance checks or deceive polygraph tests point to the defense and intelligence background of Forcepoint. Other available risk scenarios for the detection of "insider trading", "market manipulation" and "conduct risk" monitor activities such as communication, personal trading activity on the web and potential vulnerabilities of high-salaried employees to coercion in the financial sector.[62] However, the use of Forcepoint's software extends far beyond military, intelligence and banking. The company's customers for its UEBA, data loss prevention, and insider threat systems include energy providers, telecommunications firms, hospitals, airlines, manufacturers, marketing agencies[63] and retailers[64] — including in the UK, Switzerland, Italy, France, Austria[65] and many other countries. Even if extensive behavioral monitoring for cybersecurity purposes was only used in high-security departments and exclusively for sensitive personnel, it significantly affects the rights and freedoms of employees and is prone to misuse by employers.

Activity data from millions of devices. Overall, Forcepoint offers a wide range of cybersecurity software, from firewalls to web and email security.[66] Its insider threat system monitors "more than 1 million endpoints" across various large corporations and government organizations, according to the company.[67] Across its cybersecurity solutions, Forcepoint claims to receive and analyze 5 billion activity records per day from 900 million devices, including "inputs from Facebook".[68]

Data sharing across systems. Forcepoint's data loss prevention (DLP) system,[69] which aims to prevent customer data, trade secrets and other sensitive corporate information from leaving the organization's IT infrastructure, can be integrated with the UEBA system. The DLP system can send data to Forcepoint's UEBA system, which then utilizes it for risk profiling and returns risk scores about employees and other entities to the DLP system.[70] As detailed in section 4.8.1, Forcepoint's firewall,[71] cloud security,[72] access security[73] and data loss prevention (DLP)[74] systems can share data with other cybersecurity systems, for example, those provided by Microsoft.[75]

---

[62] https://www.ctc-g.com.sg/wp-content/uploads/2021/12/Forcepoint-UEBA-User-Entity-Behavior-Analytics.pdf [24.2.2024]
[63] https://www.forcepoint.com/customer-stories/communisis [26.2.2024]
[64] https://www.forcepoint.com/sites/default/files/resources/files/brochure_sureview_insider_threat_en.pdf [26.2.2024]
[65] https://www.forcepoint.com/customer-stories/enterprise-uk-hospital, https://www.forcepoint.com/customer-stories/chelsea-westminster-hospital, https://www.forcepoint.com/customer-stories/communisis, https://www.forcepoint.com/customer-stories/evalueserve, https://www.forcepoint.com/customer-stories/private-swiss-bank, https://www.forcepoint.com/customer-stories/danieli-officine-meccaniche, https://www.forcepoint.com/customer-stories/toyota-motor-italia, https://www.forcepoint.com/customer-stories/italian-insurer, https://www.forcepoint.com/customer-stories/paprec-group, https://www.forcepoint.com/customer-stories/gg-group [26.2.2024]
[66] https://www.forcepoint.com/products [5.2.2024]
[67] https://www.forcepoint.com/sites/default/files/resources/files/brochure_forcepoint_insider_threat_en.pdf [26.2.2024]
[68] Forcepoint (2018): Forcepoint Web Security, https://www.content.shi.com/SHIcom/ContentAttachmentImages/SharedResources/FBLP/Forcepoint/brochure_forcepoint_web_security_en.pdf [26.2.2024]
[69] https://www.forcepoint.com/product/dlp-data-loss-prevention [26.2.2024]
[70] p. 47, https://www.websense.com/content/support/library/ueba/v33/user_manual/user_manual.pdf [15.2.2024]
[71] https://learn.microsoft.com/en-us/azure/sentinel/data-connectors/recommended-forcepoint-ngfw-via-ama [20.3.2024]
[72] https://learn.microsoft.com/en-us/azure/sentinel/data-connectors/recommended-forcepoint-csg-via-ama [20.3.2024]
[73] https://learn.microsoft.com/en-us/azure/sentinel/data-connectors/recommended-forcepoint-casb-via-ama [20.3.2024]
[74] https://learn.microsoft.com/en-us/azure/sentinel/data-connectors/forcepoint-dlp [20.3.2024]
[75] https://learn.microsoft.com/en-us/azure/sentinel/data-connectors/recommended-forcepoint-casb-via-ama [20.3.2024]

3.7 Employee privacy and data protection?

The names of employees displayed in Forcepoint's UEBA system can be optionally replaced with pseudonyms, but the original names can still be accessed if required.[76] A number of different roles provide access to different functionality and data.[77] Forcepoint states that the system is "designed" to provide customers with "configuration options that will support their efforts to ensure" their use of the system is "GDPR compliant".[78] Employers can customize the types of data ingested into the system[79] and the behavioral risk models for profiling and scoring.[80] In a German-language presentation from 2017, Forcepoint promised "full GDPR compliance" and emphasized that the system would not put employees under general suspicion. Employers could implement the principle of multi-party verification and comply with "all requirements" of a German works council.[81] A comprehensive assessment of these measures regarding the GDPR and labor law in Germany and other European countries is beyond the scope of this study.

[76] Forcepoint (2021): Forcepoint Behavioral Analytics, USER GUIDE FOR VERSION 3.4.1, 17.9.2021, p. 36ff, https://www.websense.com/content/support/library/ueba/v341/user_guide/FBA_User_Guide_3.4.1.pdf [26.2.2024]

[77] Ibid., p. 43 and p. 36ff

[78] Forcepoint (2020): Forcepoint Behavioral Analytics, Management of Personal Data, https://www.websense.com/content/support/library/ueba/v332/GDPR/manage_personal_data.pdf [26.2.2024]

[79] Ibid.

[80] p. 48, https://www.websense.com/content/support/library/ueba/v341/user_guide/FBA_User_Guide_3.4.1.pdf [26.2.2024]

[81] Forcepoint (2017), in German "vollständige Konformität zu DSGVO/GDPR", "kein[en] wahllosen Generalverdacht", „Mehr-Augen-Prinzip", „Einhaltung aller Anforderungen eines [Betriebsrats]" the PDF document is not available online anymore: https://www.it-sa.de/CDB/download/26538302-2793-4f06-8150-044cd33c7d54 [18.3.2021]

4. Microsoft's cybersecurity and risk profiling systems Sentinel and Purview

Microsoft offers a wide range of software products for cybersecurity and risk profiling that are deeply integrated with its cloud-based Microsoft 365 system and that can be integrated with other enterprise software.[82]

Microsoft Defender promises to protect a company's Microsoft 365 system, user accounts, employee devices, mailboxes and other cloud-based applications from both malware and advanced cyberattacks.[83] **Microsoft Sentinel** is a SIEM system[84] that collects and analyzes data from "users, devices, applications and infrastructure" in order to detect and investigate cybersecurity threats and incidents.[85] Sentinel includes **UEBA functionality**[86] that calculates risk scores and promises to detect "anomalous" employee behavior based on log data and ongoing profiling.[87] Microsoft's UEBA technology can also be utilized in other systems such as Defender.[88] Furthermore, Microsoft offers systems that focus on the detection of "insider threats" and on data loss prevention (DLP). Both are part of **Microsoft Purview**,[89] which also provides "risk and compliance" functionality to monitor employee communication ("communication compliance")[90] and to search for information associated with employees ("eDiscovery").[91] As shown in Figure 6 (right), Microsoft also considers its identity and access management system Entra, formerly known as Azure Active Directory,[92] and its device management systems[93] as part of its security solutions.

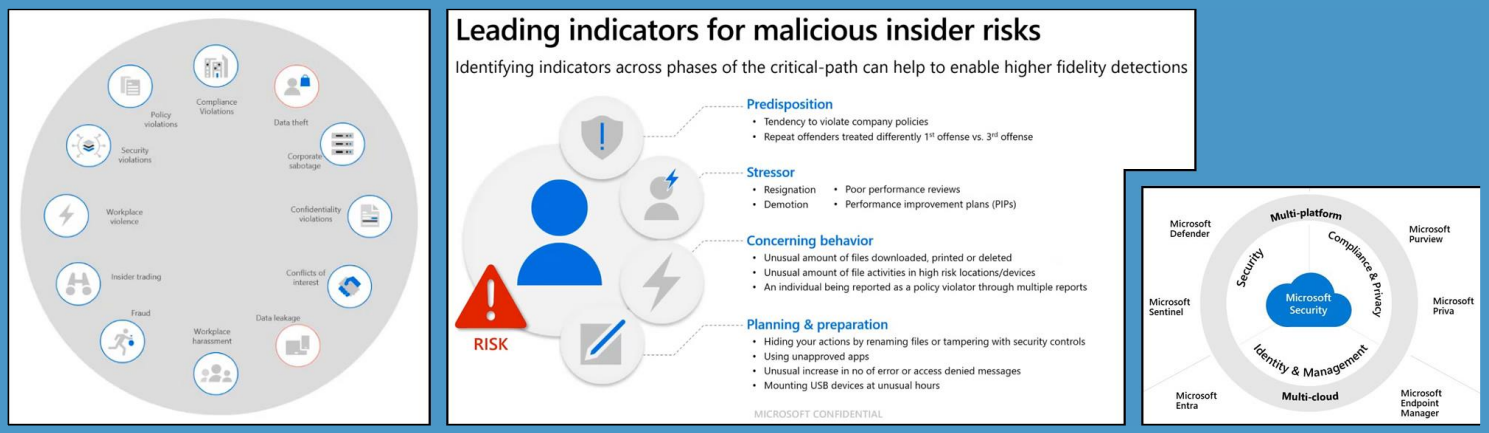

Figure 6: Identifying different types of "insider risks" and overview of security products (Microsoft)[94]

[82] https://learn.microsoft.com/en-us/microsoft-365/security [6.3.2024]
[83] https://learn.microsoft.com/en-us/microsoft-365/security/defender/microsoft-365-defender, https://learn.microsoft.com/en-us/microsoft-365/security/defender/criteria, https://learn.microsoft.com/en-us/defender-cloud-apps/editions-cloud-app-security-o365 [6.3.2024]
[84] See section 2.1
[85] https://learn.microsoft.com/en-us/azure/sentinel/overview [6.3.2024]
[86] See section 2.1
[87] https://learn.microsoft.com/en-us/azure/sentinel/identify-threats-with-entity-behavior-analytics [6.3.2024]
[88] https://learn.microsoft.com/en-us/defender-cloud-apps/tutorial-ueba [6.3.2024]
[89] https://learn.microsoft.com/en-us/purview/ [6.3.2024]
[90] https://learn.microsoft.com/en-us/purview/communication-compliance [6.3.2024]
[91] https://learn.microsoft.com/en-us/purview/ediscovery [6.3.2024]
[92] https://learn.microsoft.com/en-us/entra/identity/ [6.3.2024]
[93] https://learn.microsoft.com/en-us/mem/ [6.3.2024]
[94] Figures © Microsoft, Graham Hosking. The figures serve as basis for the discussion of the corporate practices examined in this study. Sources: video min 1:37 and 3:54 https://www.youtube.com/watch?v=Q46lAPdUx1I, https://www.microsoft.com/en-us/security/blog/2022/07/19/how-microsoft-security-partners-are-helping-customers-do-more-with-less/ [19.10.2023]

4.1 Insider risk management with Microsoft Purview

Microsoft's insider risk management system, which is part of the company's Purview product, promises to help organizations "detect, investigate, and act on malicious and inadvertent activities".[95] The system addresses a broad range of "**illegal, inappropriate, unauthorized, or unethical behavior** and actions" by employees[96] including security, policy and compliance violations, fraud, data theft, corporate sabotage, insider trading, conflicts of interest and workplace harassment (Figure 6, left). To identify employees who may become "insider risks", Microsoft suggests focusing on those who have a "**predisposition**" or "tendency" to "violate company policies", were exposed to "stressors", such as "resignation", "demotion", "poor performance reviews" or being put on "performance improvement plans", and who then show "concerning behavior" like "unusual" file or app activity (Figure 6, center).

4.1.1 Aggregate analysis of employee activity data

Organizations can initially activate Purview's **analytics functionality** to "quickly identify potential risk areas" across the organization without much configuration.[97] If activated, the system analyzes up to ten days of activity data from Microsoft 365 and Entra. Optionally, it can also access data from the HR system.[98] Figure 7 (right) shows how Microsoft displays **aggregate results** from a "scan" that involved the analysis of activity data on 23,000 employees in an organization. The system suggests that 1.3% of employees performed "potential data leak activities" and 5.9% of employees "with a resignation date" performed "potential data theft activities".

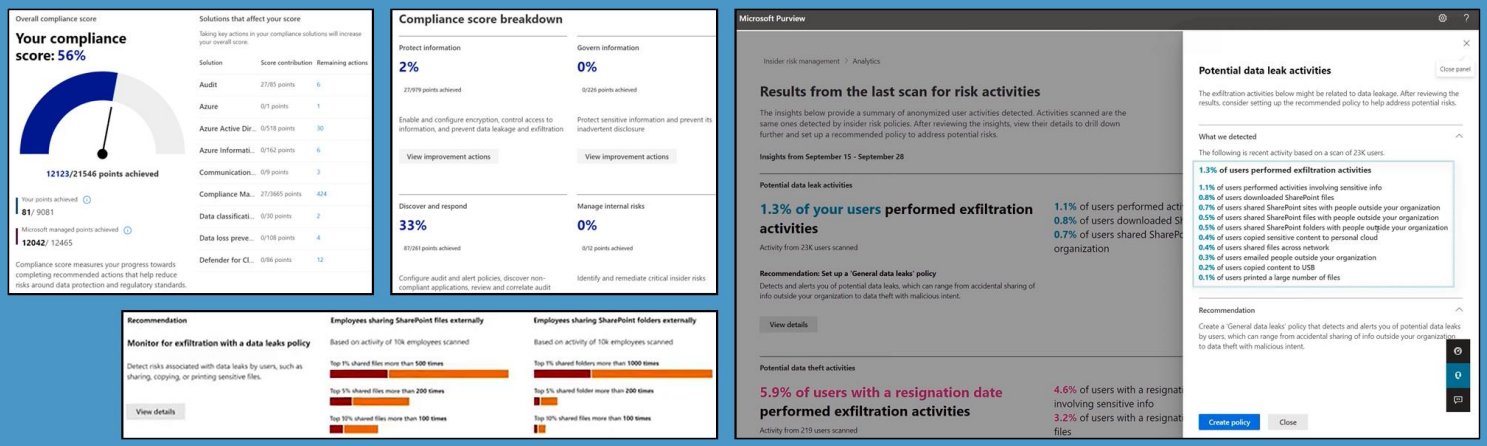

Figure 7: "Compliance scores" and assessing insider risks based on employee activity data (Microsoft Purview)[99]

Focusing on specific behaviors, 1.1% of employees performed activities "involving sensitive info", 0.5% shared files with people outside the organization, 0.2% copied content to USB devices and 0.1% printed a large number of files. Thus, Microsoft **recommends** setting up a "data leaks" policy that continuously analyzes and profiles employee activity. Figure 7 (bottom right) shows another example report on 10,000 employees that displays details on

[95] https://learn.microsoft.com/en-us/purview/insider-risk-management [7.3.2024]
[96] Ibid.
[97] https://learn.microsoft.com/en-us/purview/insider-risk-management-settings-analytics [7.3.2024]
[98] Ibid.
[99] Figures © Microsoft. The figures serve as basis for the discussion of the corporate practices examined in this study. Sources: https://learn.microsoft.com/en-us/purview/compliance-manager-setup, video min 5:59 https://www.youtube.com/watch?v=Ynkfu8OF0wQ, video min 0:58 https://www.youtube.com/watch?v=5c0P5MCXNXk [19.10.2023]

how often they shared files externally and recommends to "monitor for exfiltration with a data leaks policy".

Incentivizing employee monitoring. More broadly, Purview promises organizations the ability to assess their degree of compliance with "policies, industry standards and regional regulations".[100] As demonstrated in Figure 7 (top left), Microsoft's "compliance manager" dashboard displays an "overall compliance score" for the entire organization, which "measures" their "progress towards completing recommended actions that help reduce risks around data protection and regulatory standards". The system awards "points" for setting up and configuring different products such as Defender, "data loss prevention" and "communication compliance". Setting up the insider risk system, for example, contributes up to 12 points to the overall compliance score (Figure 7, top center). As such, Microsoft uses quantification and game mechanisms to systematically incentivize employers to set up and configure several security and compliance products, some of them involving extensive employee monitoring.

4.1.2 Profiling employee behavior based on activity data and "risk policies"

Microsoft's insider risk system offers different **risk policies** that "define the types of risks to identify and detect".[101] Depending on which risk policies are activated, the system monitors certain types of activities and starts assigning **risk scores** to employees after detecting a "triggering event" such as a "resignation", "poor performance", a "job level change" or the "evasion of security controls", as indicated in Figure 8 (left). Monitored employee activities may include everything from logging in and creating, downloading, copying, uploading, printing and deleting files to accepting meeting invitations and sending emails or chat messages. It can also include movements in physical space, such as accessing buildings or badging into conferencing rooms (Figure 8, left).

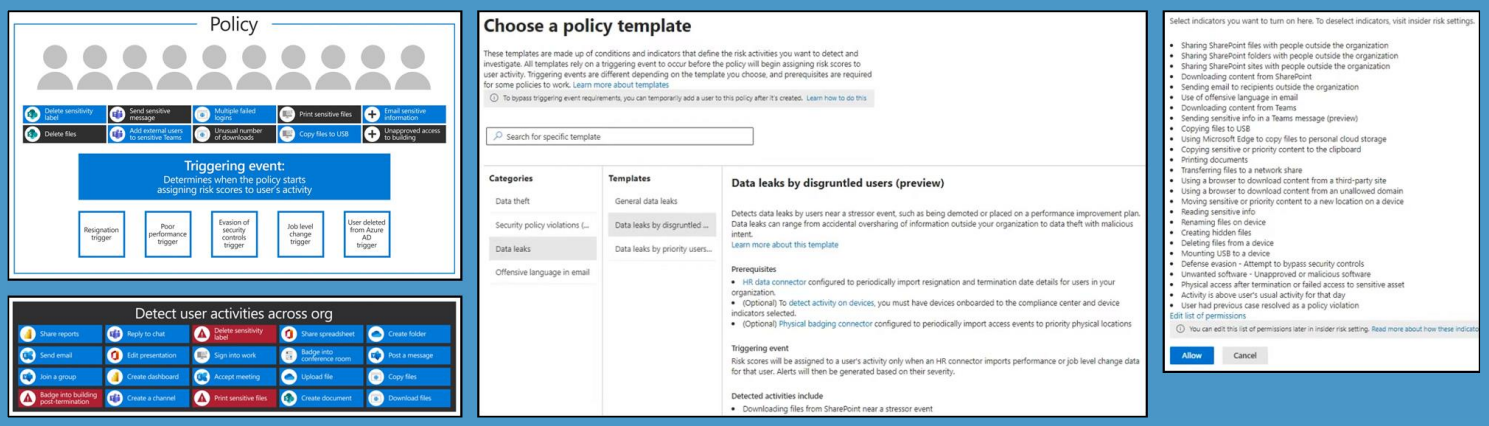

Figure 8: Insider risk policies and employee activity monitoring (Microsoft Purview)[102]

Microsoft provides a number of pre-configured **risk policy templates** that address data theft, data leaks, security policy violations and "risky browser usage". Each policy monitors certain activities and may have certain prerequisites.[103] Figure 8 (center) shows how a policy addressing "data leaks" is being added to the system. It promises to detect data leaks by **"disgruntled" employees** who are "near a stressor event, such as being demoted or placed on

[100] https://learn.microsoft.com/en-us/purview/purview-compliance [7.3.2024]
[101] https://learn.microsoft.com/en-us/purview/insider-risk-management [7.3.2024]
[102] Figures © Microsoft. The figures serve as basis for the discussion of the corporate practices examined in this study. Sources: video min 1:47 and 0:23 https://learn.microsoft.com/en-us/purview/insider-risk-management-solution-overview, https://www.microsoft.com/en-us/videoplayer/embed/RE4OUXB; video min 1:58 and 3:58 https://www.youtube.com/watch?v=kudK5ajZTUo [19.10.2023]
[103] https://learn.microsoft.com/en-us/purview/insider-risk-management-policy-templates [7.3.2024]

a performance improvement plan". In order to access HR information, it requires the "HR data connector" to be set up. Optionally, it can access activity data from employee devices and physical badging data. The system starts to calculate risk scores for employees after it has detected relevant changes in performance or job level data.

Figure 8 (right) shows the types of activities that can be monitored as part of a risk policy. Microsoft refers to these activities as **risk indicators**. In this example, monitored employee behaviors include file, printing, browser, email and clipboard activity. In addition, this risk policy tracks physical access to "sensitive assets", attempts to "bypass security controls", the use of unapproved software and the contents of email and Teams communications. The policy even includes the "use of offensive language in email" (see section 4.2). It also generally assesses if someone's "activity is above user's usual activity for that day" and if an employee conducted a "policy violation" in the past.

Disgruntled employees, risky web browsing and patient data misuse. Risk policies such as "data leaks by risky users" and "security policy violations by risky users" specifically target employees who may become risks because they "experience employment stressors" such as "performance improvement notifications, poor performance reviews, changes to job level status, or email and other messages that might signal risk activities". These policies require "disgruntlement indicators" and/or a "dedicated disgruntlement policy" to be configured, according to Microsoft. The "patient data misuse" policy addresses "misuse of patient data, either by lack of awareness, negligence, or fraud by users" including "data leaks to unauthorized persons, fraudulent modification of patient records, or the theft of patient healthcare records". The "risky browser usage" policy promises to identify visits to "potentially inappropriate or unacceptable" websites. Microsoft explains that employees who "inadvertently or purposefully visit these types of websites" may expose the employer to "legal actions", "violate regulatory requirements", "elevate network security risks" or "jeopardize current and future business operations and opportunities".[104]

As detailed in section 4.7, policies can utilize different data sources. To analyze **device activity**, the employees' Windows or macOS computers must be "onboarded".[105] To analyze **web browsing activity**, organizations must install the "Microsoft Compliance Extension" for the Edge or Chrome browser on employee devices.[106]

4.1.3 Monitoring employees over time, ranking them by risk and singling them out

Based on the activated risk policies and continuous behavioral monitoring, Microsoft's insider risk management system calculates scores for incidents and employees who are considered risks. Security, compliance and risk analysts receive **alerts about "policy matches"**. Figure 9 (left) shows two low-severity alerts about potential "data leak" incidents, both detected five days ago and involving a particular named employee. In this example, the system also displays a list of four named employees whose "risk level" was assessed as "high".

The example report in Figure 9 (top center) displays two alerts matching "data theft" and "data leak" policies with medium and high severity, both involving a particular employee. The table in Figure 9 (bottom center) gives an overview of active risk policies including information about the number of suspicious employees who are "in scope" of the policy and the number of alerts. In this example, the system utilizes risk policies such as "departing employee theft", "disgruntled users security violations", "disgruntled users data leak" and "offensive language". The report in

[104] https://learn.microsoft.com/en-us/purview/insider-risk-management-policy-templates [7.3.2024]
[105] https://learn.microsoft.com/en-us/purview/insider-risk-management-settings-policy-indicators [7.3.2024]
[106] https://learn.microsoft.com/en-us/purview/insider-risk-management-browser-support [7.3.2024]

Figure 9 (right) shows a chart that tracks the number of alerts with low, medium and high severity over several days. It also displays a list of alerts that includes a numeric risk score for each alert.

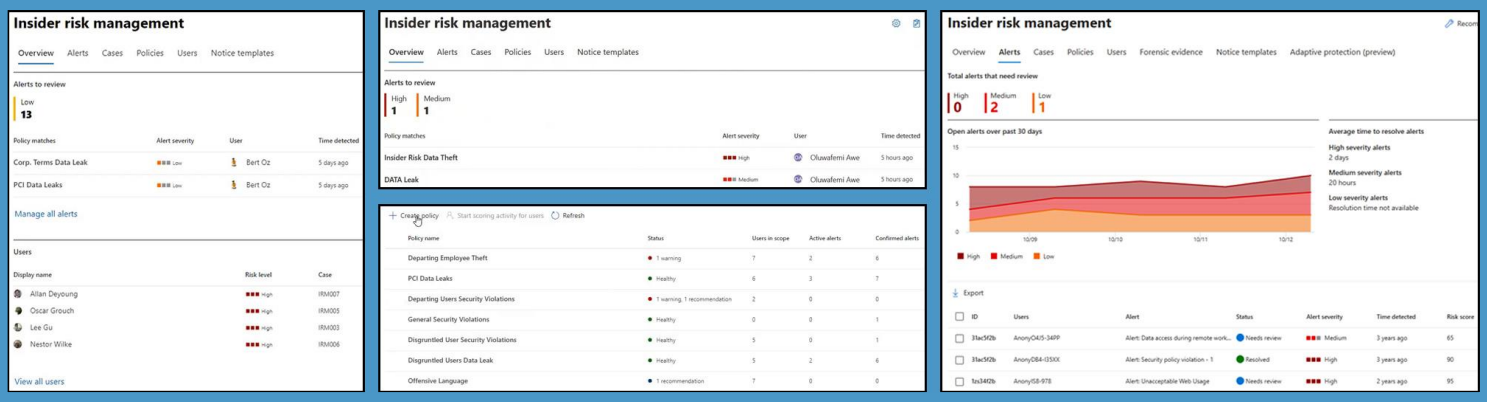

Figure 9: Profiling and ranking employees by risk level (Microsoft Purview)[107]

Detecting "sequences" and "unusual levels" of activities. Because risky activities "may not occur as isolated events" but are "frequently part of a larger sequence of events", the system can detect sequences of "two or more potentially risky activities performed one after the other that might suggest an elevated risk", according to Microsoft. It can also detect "unusual levels of risk activities" by comparing an employee's behavior with the behavior of their peers over the past 30 days based on machine learning models, for example, to identify "cumulative exfiltration" activities, i.e. employees slowly gathering and sharing confidential data over time.[108]

The report in Figure 9 (right) displays alerts that concern "data access during remote work" and "unacceptable web usage". In contrast to other example reports, this report does not display employee names but pseudonyms such as "AnonyIS8-978". The pseudonymization functionality in Microsoft Purview is further addressed in section 4.10. Figure 11 (bottom left) shows how pseudonymization can be turned on or off in the user interface.

4.1.4 Investigating past employee activity in detail

Incidents or employees who were singled out as risks can turn into **cases** for further investigation. Figure 10 (left) shows how the system displays a **timeline** of 25 "risky activities" performed by a certain employee in the past. These activities are related to an "IP theft case", i.e. the employee is suspected of stealing the employer's intellectual property. Over several months, the system detected **suspicious behavior** associated with copying, sharing and printing files, sending emails to external recipients, using "offensive language" and other activities. The person was subject of a past "insider risk case" concerning a "possible HR violation", which was confirmed as a "policy violation". Recently, a "resignation date" was set. Overall, the system assesses the case's **risk level** as 91, with 100 being the maximum. The employee is suspected of a violation of their "confidentiality obligation during departure".

The report in Figure 10 (right) shows an even more detailed timeline of activities carried out by another employee, who is also suspected of "potential IP theft". The case's **risk score** is assessed as "25" and thus as only "low". The

[107] Figures © Microsoft, Graham Hosking, Ha-Shem. The figures serve as basis for the discussion of the corporate practices examined in this study. Sources: video min 6:26 https://www.youtube.com/watch?v=Q46lAPdUx1I, video min 22:56 https://www.youtube.com/watch?v=rJlMUhrVcoU, video min 21:17 https://www.youtube.com/watch?v=Q46lAPdUx1I, https://learn.microsoft.com/en-us/purview/insider-risk-management [19.10.2023]
[108] https://learn.microsoft.com/en-us/purview/insider-risk-management-policies [7.3.2024]

system accuses the employee of "cumulative exfiltration activities" over several months, based on the detection of **unusual behaviors**. The person carried out some activities 20% more often than other employees, other activities were carried out 90% more often. According to the "activity explorer" report, the suspicious employee accessed "healthcare records" and browsed "hacking websites", "keylogger websites", "criminal activity websites", "gambling websites" and "cult websites". The list includes specific website addresses visited by the person.

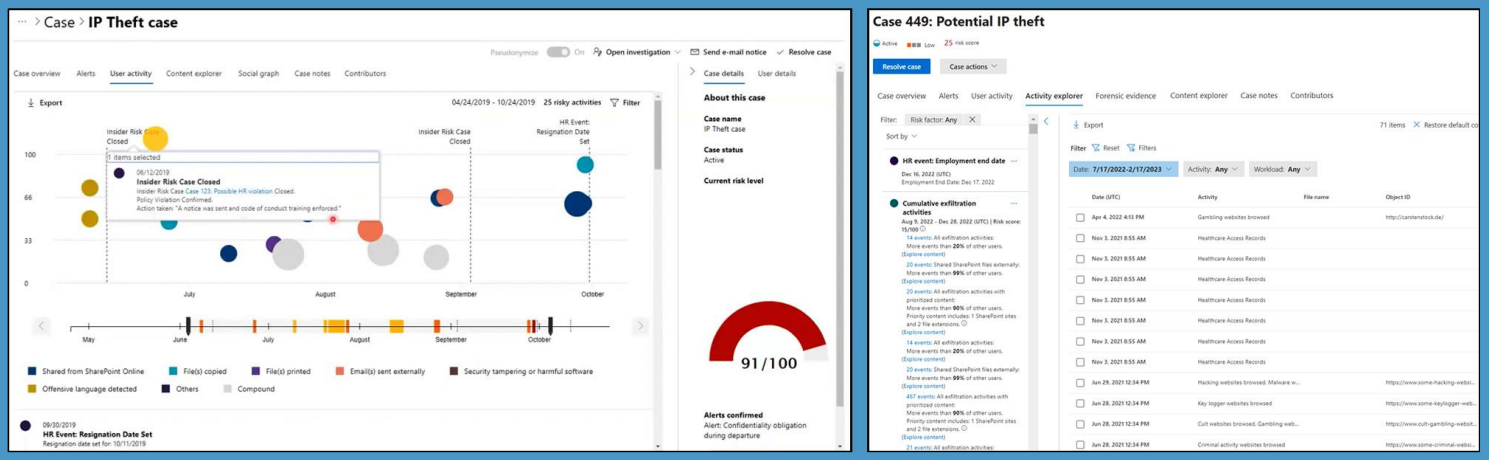

Figure 10: Investigating past employee activity (Microsoft Purview)[109]

Another example report shown in Figure 11 (top left) displays information about a named employee who is suspected of "departing employee theft". The system detected an "unusual volume of sensitive files read". Subsequently, the person moved "sensitive files" to another location and later deleted files from their Windows 10 computer.

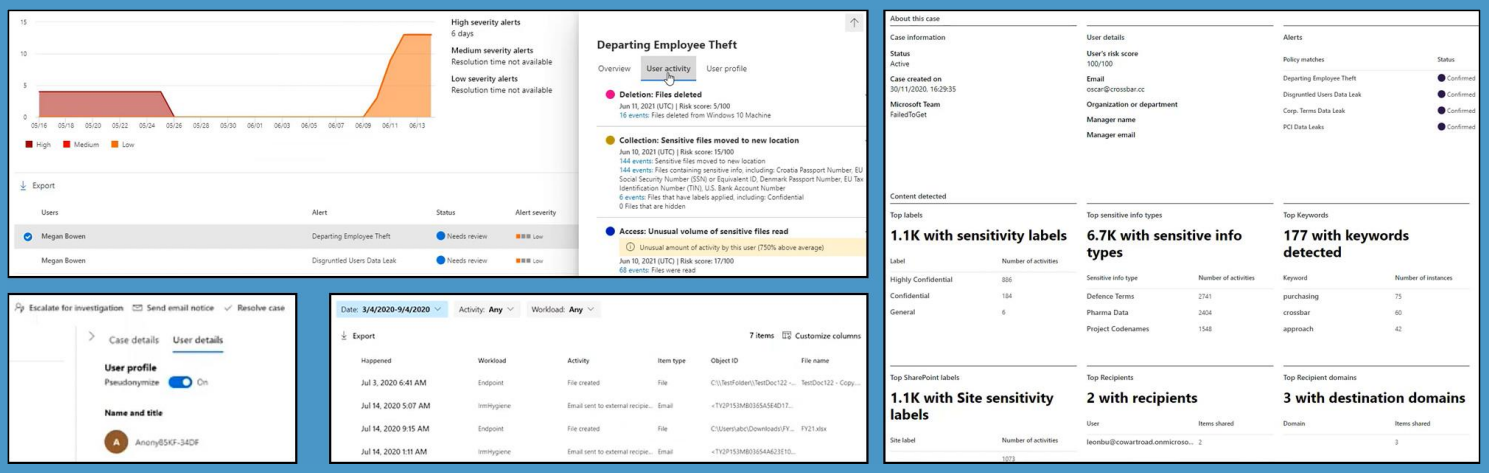

Figure 11: Investigating past employee activity (Microsoft Purview)[110]

[109] Figures © Microsoft, New Era. The figures serve as basis for the discussion of the corporate practices examined in this study. Sources: video min 39:20 https://www.youtube.com/watch?v=uTQ0dg6mlzs, https://learn.microsoft.com/en-us/purview/insider-risk-management-activities [19.10.2023]

[110] Figures © Microsoft, Graham Hosking. The figures serve as basis for the discussion of the corporate practices examined in this study. Sources: video min 11:24 and 14:00 https://www.youtube.com/watch?v=Q46lAPdUx1I, video min 16:24 https://www.youtube.com/watch?v=4Q7q1ELcI48, video min 20:54 https://www.youtube.com/watch?v=gqtXp2b36yQ [19.10.2023]

Figure 11 (bottom center) shows a report that displays information about how an employee created a particular file, including the filename, based on data recorded from the employee's computer. In addition, the person sent emails to external recipients, according to the list of activities. The lists in Figure 12 (left and bottom) also demonstrate how the system displays detailed information about **file activities** carried out by a particular named employee, from creating and modifying files to copying them to the cloud and to the clipboard.

As shown in Figure 10 (right), the system can also display information about **communication contents and recipients**. This report concerns a case associated with a named employee whose activities triggered several risk policies addressing "employee theft" and data leaks by "disgruntled" employees. The system assigned the person the maximum risk score of 100 and detected a large number of activities associated with "confidential" and "highly confidential" content. Many activities concerned sensitive information labeled "defence terms" and "pharma data". The report also lists the detected keywords, the email addresses of email recipients and top "destination domains", which indicates that the insider risk system also analyzes data on persons outside the organization.

Microsoft's insider risk system provides access to the **full contents** of documents and communication activities associated with employees who were assessed as risks via the "content explorer". It can display the contents of files, emails and messages from Microsoft systems such as SharePoint, Exchange, Teams and OneDrive, which can be filtered and searched by date, document author, message sender and recipients and other criteria.[111]

4.1.5 Monitoring device activities via screen recording to gather "forensic evidence"

Purview's "forensic evidence" functionality can record in detail how employees use their computer and provides access to fine-grained user interaction logs and visual representations of their screen activity.[112]

Figure 12 (right) demonstrates how a "captured clip" about activities carried out by a particular employee at a certain date and time can be accessed. The video shows the employee's Windows desktop including a file explorer window. The log on the right provides information about every user interaction over the time period of one minute.

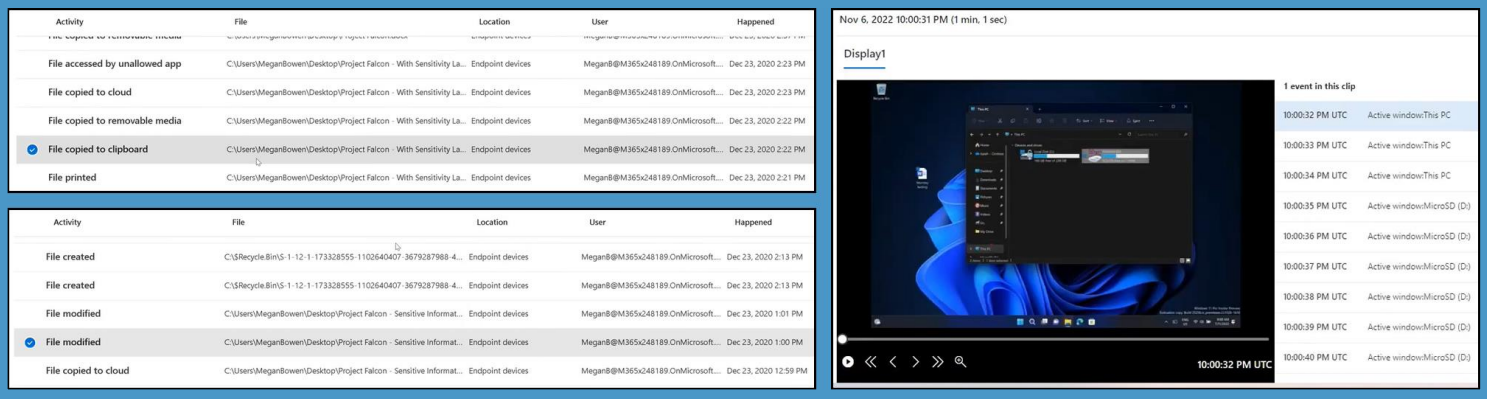

Figure 12: Investigating past employee activity and screen recording for forensic evidence (Microsoft Purview)[113]

[111] https://learn.microsoft.com/en-us/purview/insider-risk-management-content-explorer [7.3.2024]
[112] https://learn.microsoft.com/en-us/purview/insider-risk-management-forensic-evidence [12.3.2024]
[113] Figures © Microsoft. The figures serve as basis for the discussion of the corporate practices examined in this study. Sources: video min 2:08 and 3:16 https://www.youtube.com/watch?v=gkXPUNiIp84, video min 12:39 https://www.youtube.com/watch?v=zv8I6l-UmlQ [19.10.2023]

The **"forensic evidence" functionality** must be explicitly activated for specific employees, who can optionally be notified about the activation. The recording can include or exclude certain desktop applications or websites. Microsoft suggests excluding personal email and social media accounts. The system can either continuously capture activities carried out by employees or provide access only to activities that occur a few minutes before until a few minutes after a certain "trigger" activity was detected. For time-sensitive evidence-gathering needs, it allows for quick activation of "continuous capturing" without much restrictions, which can help employers to prevent a "potentially risky activity from being missed or going undetected", according to Microsoft.[114]

4.2 Monitoring employee conversations with Purview "communication compliance"

The "communication compliance" system, which is also part of Microsoft's broader data security, risk and compliance system Purview, allows companies to automatically analyze and scan everything employees write or say, both to each other and to persons outside the organization. It aims to "minimize communication risks" by helping employers **"detect, capture, and act on potentially inappropriate messages"** in emails, chats, documents, images, meeting transcripts and other contents. The system can monitor and analyze communication activities in Microsoft 365, Exchange, Teams, Viva Engage and Copilot.[115] It can also import data from custom data sources (e.g. databases) and third-party applications (e.g. Zoom, Slack, Cisco Webex, WhatsApp, Signal), and it can even monitor text messages and voice calls from standard mobile phones,[116] as detailed in section 4.7.

Microsoft refers to the communication compliance system both as a "compliance" and an "insider risk" solution.[117] It allows for extensive monitoring of communication for **very different purposes** ranging from the detection of "profanity", "offensive language" and "inappropriate text" to the detection of corporate sabotage, money laundering, bribery, insider trading and unauthorized disclosure of confidential information.[118] Communication and file content that was assessed as inappropriate raises an alert and can then be further investigated. Optionally, the system can notify the associated employee or directly remove inappropriate content. For example, it can **automatically block messages** in Microsoft Teams and instead display a notification that explains that the message was removed because of a "policy violation".[119] According to Microsoft, organizations can use the system to "identify and manage potential legal exposure and risk" and ensure that their employees comply with "regulatory compliance standards" and with "acceptable use, ethical standards, and other corporate policies in all their business-related communications".[120]

Microsoft incentivizes employers to enable the communication compliance system. As described in section 4.1.1 and shown in Figure 7 (top left), setting up communications monitoring contributes up to 9 points to the organization's overall "compliance score" in the "compliance manager" dashboard. While employers can specifically include or exclude certain groups of employees from communication monitoring, Microsoft recommends monitoring

[114] https://learn.microsoft.com/en-us/purview/insider-risk-management-forensic-evidence [12.3.2024]
[115] https://learn.microsoft.com/en-us/purview/communication-compliance [13.3.2024]
[116] https://learn.microsoft.com/en-us/purview/archive-third-party-data [13.4.2024]
[117] https://learn.microsoft.com/en-us/purview/communication-compliance [13.3.2024]
[118] https://learn.microsoft.com/en-us/purview/communication-compliance-policies, https://learn.microsoft.com/en-us/purview/trainable-classifiers-definitions [13.3.2024]
[119] https://learn.microsoft.com/en-us/purview/communication-compliance [13.4.2024]
[120] Ibid.

"all" employees in an organization at least for "harassment or discrimination detection".[121] The system can also analyze all employee communication in order to make recommendations about which kind of monitoring should be set up.[122] Figure 13 (left) shows how an organization is told that their employees "recently sent 120 emails that might contain inappropriate content". The report further specifies that 27 messages "might" contain "profanity", 32 might contain "targeted harassment" and 47 messages might contain a "threat". Consequently, Microsoft recommends to "start monitoring communications" to detect "inappropriate" content "now".

4.2.1 Detecting suspicious communications content with "policies" and AI-based "classifiers"

Similar to the insider risk system, Microsoft's communication compliance system uses different "policies" to detect certain types of communications that are assessed as risks. The example dashboard in Figure 13 (left) demonstrates how the system detected 25 emails, chat messages or other content items that match the "profanities" policy, three items that match the "inappropriate images" policy and 224 items that match the "insiders" policy.

Policy alerts. Figure 13 (second from the left) shows another example report that displays a chart about recent matches to the "regulatory compliance" and "offensive and threatening language" policies over several days. In addition, the report displays a ranking of named employees "with the most policy matches". The system also provides a report that displays a list of employees with policy matches, including information about whether the suspicious communication activity was reviewed, whether the reviewer assessed it as compliant or non-compliant and whether the employee was notified about the issue (Figure 13, bottom right).

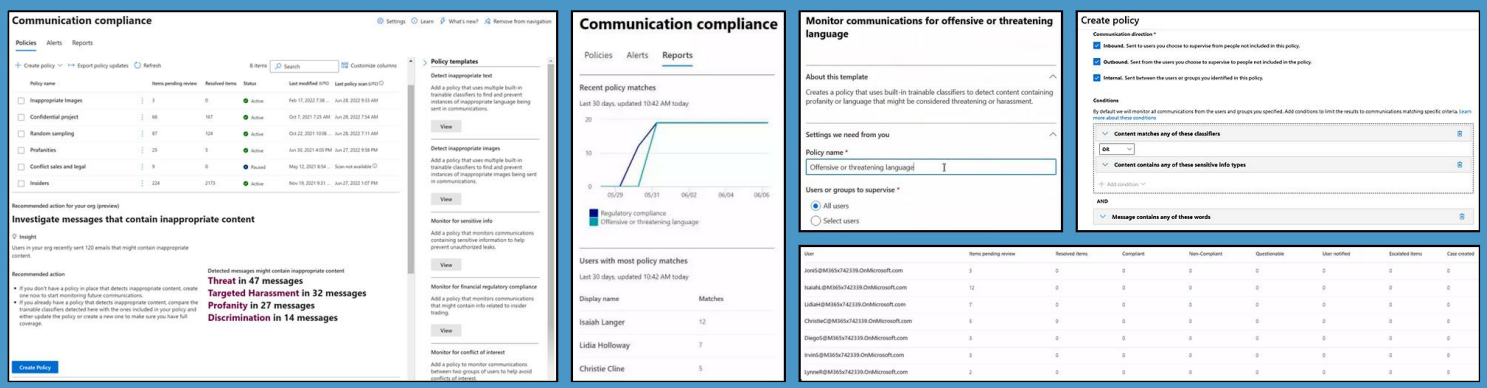

Figure 13: Detecting suspicious communications content with "policies" (Microsoft Purview)[123]

A simple communication compliance policy can monitor messages and other contents for certain **keywords and phrases**. The system raises an alert if these keywords are detected. A policy can also use **classifiers** that promise to recognize certain types of communication content such as "offensive language" or "corporate sabotage" based on machine learning and AI models. Figure 13 (top right) shows the creation of a new policy, which also allows the

[121] "Consider adding all users in your organization as in-scope for your communication compliance policies". Identifying specific users as in-scope for individual policies are useful in some circumstances, however most organizations should include all users in communication compliance policies optimized for harassment or discrimination detection", https://learn.microsoft.com/en-us/purview/communication-compliance-plan [13.3.2024]

[122] https://learn.microsoft.com/en-us/purview/communication-compliance-policies [13.3.2024]

[123] Figures © Microsoft, Joanne Klein, Quest, Ha-Shem. The figures serve as basis for the discussion of the corporate practices examined in this study. Sources: https://www.linkedin.com/pulse/purview-communications-compliance-101-beau-faull, https://www.slideshare.net/joanneklein/communication-compliance-in-microsoft-365, video min 35:33 https://www.youtube.com/watch?v=rJlMUhrVcoU, https://practical365.com/how-to-create-a-communication-compliance-policy/, https://www.slideshare.net/joanneklein/communication-compliance-in-microsoft-365 [19.10.2023]

option to choose whether the new policy should monitor internal communication between employees, "inbound" communications sent from external persons to employees or "outbound" communications sent from employees to external persons. Microsoft provides a number of pre-built **policy templates**, some of which are listed in Table 3:

Area	Policy template	Data sources	Direction	Default conditions and pre-trained classifiers
Inappropriate content	Detect inappropriate content	Teams	Inbound, Outbound, Internal	Hate, Violence, Sexual, Self-harm classifiers
Inappropriate text	Detect inappropriate text	Exchange, Teams, Viva Engage	Inbound, Outbound, Internal	Threat, Discrimination, and Targeted harassment classifiers
Inappropriate images	Detect inappropriate images	Exchange, Teams, Viva Engage	Outbound, Internal	Adult and Racy image classifiers
Copilot interactions	Detect Copilot interactions	Copilot for Microsoft 365	Inbound, Outbound, Internal	-
Sensitive information	Detect sensitive info types	Exchange, Teams, Viva Engage	Inbound, Outbound, Internal	Sensitive information, out-of-the-box content patterns and types, custom dictionary option
Regulatory compliance	Detect financial regulatory compliance	Exchange, Teams, Viva Engage	Inbound, Outbound	Customer complaints, Gifts & entertainment, Money laundering, Regulatory collusion, Stock manipulation, and Unauthorized disclosure classifiers
Conflict of interest	Detect conflict of interest	Exchange, Teams, Viva Engage	Internal	-

Table 3: "Communication compliance" policy templates (Microsoft Purview)[124]

The "inappropriate text" policy template, for example, uses a number of classifiers that promise to detect harassment, discrimination and threats in inbound, outbound and internal communication carried out via Microsoft 365, Exchange, Teams and Viva Engage. The "regulatory compliance" policy template for financial organizations uses classifiers that promise to detect employee communication that indicates money laundering, stock manipulation, collusion, bribery, customer complaints and unauthorized disclosure of confidential information.

AI-based content classifiers. As of 2024, Microsoft offers 59 "pre-trained" content classifiers that promise to detect certain types of communication,[125] some of which are listed in Table 4. Microsoft refers to them as "pre-trained" classifiers because they are provided out of the box. Organizations can also "train" their own content classifiers by feeding the system with at least 50 text samples that "strongly represent the type of content" they want to detect. Subsequently, a machine learning model is created.[126] A communication compliance policy can then use one or several classifiers – either pre-trained by Microsoft or created by employers – to monitor communication.

As Table 4 shows, the pre-trained classifiers provided by Microsoft serve very different purposes. The "corporate sabotage" classifier, for example, "detects messages that may mention acts to damage or destroy corporate assets or property". Other classifiers promise to detect messages that point to other types of employee misconduct, from sharing confidential information ("unauthorized disclosure" classifier) to bribery in the form of "exchanging gifts or entertainment in return for service" ("gifts & entertainment" classifier). The "money laundering", "stock manipulation" and "regulatory collusion" classifiers detect message content that indicate violations of regulations in the financial industry and other sectors. Announced in 2022, the "workplace collusion" classifier was described as a means to detect "secretive actions such as concealing information or covering instances of a private conversation,

[124] https://learn.microsoft.com/en-us/purview/communication-compliance-policies [13.3.2024]
[125] https://learn.microsoft.com/en-us/purview/trainable-classifiers-definitions [13.3.2024]
[126] https://learn.microsoft.com/en-us/purview/trainable-classifiers-get-started-with [12.3.2024]

interaction, or information".[127] Yet other classifiers aim to detect the use of "offensive language", "harassment", "discrimination" or "threats to commit violence or do physical harm or damage to a person or property". The "adult, racy, and gory images" classifier promises to detect "inappropriate" images in emails and chats.

Content classifier	Description (original quotes)
Corporate sabotage	Detects messages that may mention acts to damage or destroy corporate assets or property
Intellectual property	Detects content in intellectual property related categories such as trade secrets and similar confidential information
Unauthorized disclosure	Detects sharing of information containing content that is explicitly designated as confidential or internal to unauthorized individuals
Money laundering	Detects signs that may suggest money laundering or engagement in acts to conceal or disguise the origin or destination of proceeds
Gifts & entertainment	Detects messages that may suggest exchanging gifts or entertainment in return for service, which violates regulations related to bribery
Regulatory collusion	Detects messages that may violate regulatory anti-collusion requirements such as an attempted concealment of sensitive information
Stock manipulation	Detects signs of possible stock manipulation, such as recommendations to buy, sell or hold stocks that may suggest an attempt to manipulate the stock price
Threat	Detects a specific category of offensive language text items related to threats to commit violence or do physical harm or damage to a person or property
Harassment	Detects a specific category of offensive language text items related to offensive conduct targeting one or multiple individuals based on the following traits: race, ethnicity, religion, national origin, gender, sexual orientation, age, disability
Discrimination	Detects explicit discriminatory language
Profanity	Detects a specific category of offensive language text items that contain expressions that embarrass most people
Adult, racy, and gory images	Detects images that are potentially inappropriate. Scanning and detection are supported for Exchange Online email messages, and Microsoft Teams channels and chats
Personal financial information	Detects documents related to different personal financial records consisting of financial statements, real estate planning and retirement plans. Consists of details of all assets and liabilities held by an individual
Employee disciplinary action	Detects files relating to disciplinary action including a reprimand or corrective action in response to employee misconduct, rule violation, or poor performance
Non-disclosure agreement	Detects nondisclosure agreements (NDAs)
Customer complaints	The customer complaints classifier detects feedback and complaints made about your organization's products or services
Meeting notes	Detects documents and notes containing information specific to meetings
Financial statement	Detects financial statements like income statement, balance sheet, cash flow statement, statement of changes in equity
Business context	Detects presence of business-related content such as organizational structure, policy updates, contracts, HR policies, crucial financial data such as revenue and profits, healthcare forms, employee contracts
Procurement	Detects content in categories of bidding, quoting, purchasing, and paying for supply of goods and services
Healthcare	Detects content in medical and healthcare administration aspects such as medical services, diagnoses, treatment, claims
Safety records	Detects documents that are related to facility/factory safety

Table 4: "Communication compliance" content classifiers (Microsoft Purview)[128]

Several pre-trained classifiers provided by Microsoft, as listed in Table 4, promise to identify certain types of content rather than making an assessment about employee behavior. The system can, for example, recognize messages and documents that contain information related to trade secrets, financial statements, contracts, procurement, meeting notes, customer complaints, healthcare diagnoses, facility safety records and personal financial information. These

[127] https://pupuweb.com/mc387035-microsoft-purview-additional-classifiers-communication-compliance-preview/ [13.3.2024]
[128] https://learn.microsoft.com/en-us/purview/trainable-classifiers-definitions [12.3.2024]

classifiers can be used in combination with other classifiers associated with employee behavior. As such, these classifiers can help organizations either to protect sensitive information or to single out employees.

The list of classifiers shows that Microsoft's communication monitoring system addresses a diverse range of purposes, from detecting offensive language to complying with legal obligations to preventing criminal conduct. While observing all communication activities may be legitimate for some purposes for certain groups of employees, it can be considered **intrusive and disproportionate** in other cases. In any case, it can have problematic side effects, from inaccurate suspicions to spying on employees. In 2022, Microsoft announced it was introducing a "leavers" classifier that promised to detect "messages that explicitly express intent to leave the organization". As this led to a heated public debate, the company stated later in the year that it would not introduce the "leavers" classifier "at this time".[129] As employers can train their own classifiers, they can, however, easily create a similar classifier by themselves.

Each communication compliancy policy can store up to 100 GB of content or 1 million messages, which suggests that the system is prepared for **corporate mass surveillance**.[130] With regards to accuracy, Microsoft emphasizes that "there are limitations to any artificial intelligence solution" including "false positive" and "false negative" detection.[131] A number of pre-trained classifiers support only English-language content; some classifiers support up to twelve languages including Spanish, Portuguese, French, German, Italian, Dutch, Japanese, Korean, Arabic and Chinese.[132] Similar to the insider risk system, Microsoft's communication compliance system allows organizations to replace employee names with pseudonyms in the user interface[133] (see also section 4.10).

4.2.2 Investigating message contents, images and meeting recordings

After the communication compliance system has created alerts about suspicious communication activities, investigators can access and review the full message contents. Figure 14 (left) shows how the system displays the content of an email that triggered the "offensive and threatening language" policy and up to five messages sent before and five sent after the suspicious email.

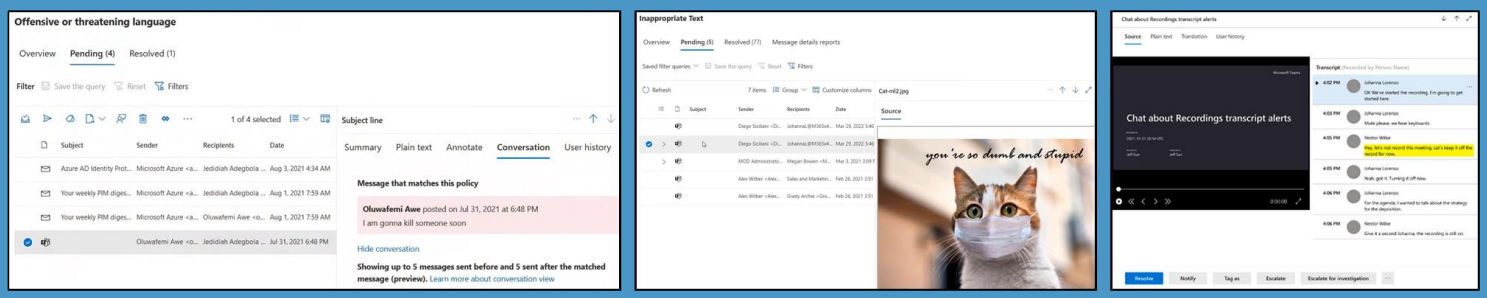

Figure 14: Investigating message contents, images and meeting recordings (Microsoft Purview)[134]

[129] https://pupuweb.com/mc387035-microsoft-purview-additional-classifiers-communication-compliance-preview/, https://www.reddit.com/r/sysadmin/comments/v3b2mn/microsoft_introducing_ways_to_detect_people/ [13.3.2024]
[130] https://learn.microsoft.com/en-us/purview/communication-compliance-policies [13.3.2024]
[131] https://learn.microsoft.com/en-us/purview/communication-compliance-solution-overview [13.3.2024]
[132] https://learn.microsoft.com/en-us/purview/trainable-classifiers-definitions [13.3.2024]
[133] https://learn.microsoft.com/en-us/purview/communication-compliance-solution-overview [13.3.2024]
[134] Figures © Microsoft, Ha-Shem. The figures serve as basis for the discussion of the corporate practices examined in this study. Sources: video min 46:29 https://www.youtube.com/watch?v=rJlMUhrVcoU, video min 12:27 https://www.youtube.com/watch?v=Ynkfu8OF0wQ, https://techcommunity.microsoft.com/t5/security-compliance-and-identity/keep-microsoft-teams-meetings-compliant-with-communication/ba-p/3933446 [19.10.2023]

The system can also analyze **text within images** via optical character recognition (OCR) technology,[135] as demonstrated in Figure 14 (center). Furthermore, it can analyze text in automated **transcripts of meetings** recorded via Microsoft Teams.[136] When a suspicious communication activity is detected in a transcript, investigators can review what participants said by reading the transcript or by watching the recorded video (Figure 14, right).[137] The system also promises to detect whether suspicious messages have a "positive", "negative" or "neutral" **sentiment**.[138]

4.3 Combining and automating insider risk detection and communication monitoring

Both the insider risk (see section 4.1) and the communication compliance system (section 4.2) are part of **Microsoft Purview**. While Microsoft refers to the former as a "data security solution" and promotes the latter as a "risk and compliance solution", both technologies are deeply intertwined.

The detection of "offensive language" via communication compliance can contribute to assessments of employees as potential "insider threats", as shown in Figure 8 (right) in section 4.1.2, Figure 9 (bottom center) in section 4.1.3 and Figure 10 (left) in section 4.1.4. Microsoft explains that "workplace stress may lead to uncharacteristic or malicious behavior" by employees that could "surface as potentially inappropriate behavior" in communication and messaging activities. Consequently, the communication compliance system can provide "risk signals" to insider risk policies, for example, by using an "inappropriate text policy", which can be automatically created when adding the pre-configured insider risk policies "data leaks by risky users" or "security policy violations by risky users".[139]

Automatically notifying managers and creating tasks. Both systems offer a wide range of functions to respond to alerts and manage cases. The insider risk system can automatically notify managers when an insider risk alert is detected for an employee. Microsoft's workflow management system "Power Automate" can be used to automate tasks, such as requesting information about a suspicious employee from an organization's human resource department or automatically creating an associated task in the task management system ServiceNow.[140] Similarly, the communication compliance system can automatically notify managers about alerts and trigger other actions via Power Automate.[141] In addition, investigators can easily create a new case[142] in Microsoft's "eDiscovery" system.[143]

Both systems can **share alert data** with Microsoft's cybersecurity systems Sentinel and Defender or with cybersecurity systems from other vendors such as Splunk,[144] a Cisco subsidiary.[145] The insider risk system can share information about user and employee risk levels with Purview's data loss prevention (DLP) system.[146]

---

[135] https://learn.microsoft.com/en-us/purview/communication-compliance [14.3.2024]
[136] https://learn.microsoft.com/en-us/purview/communication-compliance-investigate-remediate [14.3.2024]
[137] https://techcommunity.microsoft.com/t5/security-compliance-and-identity/keep-microsoft-teams-meetings-compliant-with-communication/ba-p/3933446 [14.3.2024]
[138] https://learn.microsoft.com/en-us/purview/communication-compliance-investigate-remediate [6.5.2024]
[139] https://learn.microsoft.com/en-us/purview/communication-compliance-policies [14.3.2024]
[140] https://learn.microsoft.com/en-us/purview/insider-risk-management-settings-power-automate [14.3.2024]
[141] https://learn.microsoft.com/en-us/purview/communication-compliance-investigate-remediate [14.3.2024]
[142] Ibid.
[143] See section 4.4
[144] https://learn.microsoft.com/en-us/purview/communication-compliance-siem, https://learn.microsoft.com/en-us/purview/insider-risk-management-settings-dlp-sync [14.3.2024]
[145] https://www.splunk.com/en_us/about-splunk/acquisitions.html [2.5.2024]
[146] https://learn.microsoft.com/en-us/purview/insider-risk-management-settings-dlp-sync [26.4.2024]

4.4 Searching for employee information and compiling dossiers on them with Purview "eDiscovery"

While the communication compliance system examined in section 4.2 promises to detect undesirable communication and document content as it is being created, Purview's **eDiscovery solutions** enable organizations to identify, discover, access and analyze all "electronic information" associated with certain employees or search terms, including documents, files, email communication and chats in Microsoft Exchange, OneDrive, SharePoint and Teams, in order to use it "as evidence in legal cases".[147] Microsoft explicitly states that investigations via eDiscovery should help organizations respond to legal cases involving "executives or other employees". This can involve "quickly finding and retaining" information in "email, documents, instant messaging conversations, and other content locations used by people in their day-to-day work tasks".[148] eDiscovery should help organizations respond to "external investigations" and "legal matters or internal investigations" by "discovering data where it lives" and identifying "persons of interest".[149]

Microsoft offers three different eDiscovery solutions. The **Content Search** tool allows organizations to "search for content across Microsoft 365 data sources". Building on it, the **"eDiscovery (Standard)"** system adds case management and allows certain mailboxes, chat messages or document repositories to be put on "legal hold"[150] in order to preserve contents "related to a specific investigation or for a specific person" and make deletion impossible.[151] The **"eDiscovery (Premium)"** system provides additional analysis functionality.[152]

When creating an eDiscovery case, organizations can **choose one or several employees** and then use "keywords, properties, and conditions to build search queries that return search results with the data that's most likely relevant to the case" using content from Exchange mailboxes, OneDrive files, Teams and other data sources. This can include information created or maintained by employees who are subject to the investigation ("custodian data") and related information about other employees ("non-custodial data").[153] The system can search and return large numbers of documents and conversations from many employees across the organization. Microsoft provides an example that shows an eDiscovery case that involves searching 7,309 mailboxes and resulted in 2,549,828 search hits in 5,815 mailboxes. As such, the system can be used for both **targeted and mass surveillance** in an organization.[154]

The retrieved information can include conversations that involve **persons outside the organization**. For example, it can include meeting recordings and transcripts in Teams, audio calls between employees and "external contacts" and "chats with guests".[155] It can also include prompts and responses in Microsoft Copilot.[156] Organizations can **export** all information retrieved via eDiscovery.[157] Putting Exchange mailboxes, Teams conversations and other

[147] https://learn.microsoft.com/en-us/purview/ediscovery [25.4.2024]
[148] https://learn.microsoft.com/en-us/purview/ediscovery-manage-legal-investigations [25.4.2024]
[149] https://learn.microsoft.com/en-us/purview/ediscovery-overview [25.4.2024]
[150] https://learn.microsoft.com/en-us/purview/ediscovery [25.4.2024]
[151] https://learn.microsoft.com/en-us/purview/ediscovery-teams-legal-hold [25.4.2024]
[152] https://learn.microsoft.com/en-us/purview/ediscovery [25.4.2024]
[153] https://learn.microsoft.com/en-us/purview/ediscovery-create-and-manage-cases [25.4.2024]
[154] https://learn.microsoft.com/en-us/purview/ediscovery-create-draft-collection [25.4.2024]
[155] https://learn.microsoft.com/en-us/purview/ediscovery-teams-investigation [25.4.2024]
[156] https://learn.microsoft.com/en-us/purview/ediscovery-search-and-delete-copilot-data [25.4.2024]
[157] https://learn.microsoft.com/en-us/purview/ediscovery [25.4.2024]

content locations on "legal hold" can be based on more or less specific search queries.[158] Legal holds can be either "infinite" or address certain data ranges.[159]

Microsoft's eDiscovery products are deeply integrated with its **insider risk system**.[160] In "situations where additional legal review is needed for the user's risk activity", as detected by Purview's insider risk system, organizations can "escalate the case for user investigation" via the eDiscovery system in order to "preserve, collect, review, analyze, and export content" relevant to "legal investigations".[161]

4.5 Data loss prevention (DLP) and other Purview functionality

Microsoft Purview offers a wide range of data security, risk profiling, compliance and data governance functionality. In addition to the insider risk, communication compliance and eDiscovery systems examined in the previous sections, this includes a **data loss prevention (DLP)** system that aims to prevent sensitive corporate information from leaving an organization's IT infrastructure,[162] whether intentionally or not. The DLP system promises to detect when employees perform activities that lead to inappropriate sharing of sensitive information with external parties, from intellectual property and trade secrets to personal information about an organization's customers, employees and other persons (e.g. credit card numbers, social security numbers, health records).[163] For this purpose, it **monitors communication and file activities** that involve information that was declared sensitive across Microsoft 365, Exchange, Teams, OneDrive and SharePoint. In addition, it can monitor activities performed in office applications (e.g. Word, Excel, PowerPoint), other Microsoft systems (e.g. Power BI) and third-party systems (via Microsoft's "Defender for Cloud Apps").[164]

Device and browser monitoring. The DLP system can monitor employees' Windows and macOS computers in order to detect activities such as copying files that contain sensitive information to the clipboard, to a network share or an USB device, uploading them to a cloud service, creating and renaming files, printing them, opening them with a "restricted" application or accessing them from an "unallowed" browser.[165] Additionally, organizations can install browser extensions on employee devices that monitor file activity in the web browser.[166]

Singling out employees based on behavioral monitoring. Similar to the other Purview risk profiling systems, the DLP system raises alerts about suspicious activities based on customizable "policies". Each alert contains information about the "violating action" (e.g. file copied to clipboard), the type of sensitive information detected (e.g. credit card data), the concerned communication item or file, the time the activity occurred and the user or employee whose activity triggered the alert.[167] As such, the system can single out employees based on monitoring their behavior. In contrast to Microsoft's insider risk system, it does not create behavioral profiles about employees over

[158] https://learn.microsoft.com/en-us/purview/ediscovery-create-holds [25.4.2024]
[159] Ibid.
[160] https://learn.microsoft.com/en-us/purview/ediscovery [25.4.2024]
[161] https://learn.microsoft.com/en-us/purview/insider-risk-management-cases [25.4.2024]
[162] https://learn.microsoft.com/en-us/purview/dlp-learn-about-dlp [26.4.2024]
[163] https://www.microsoft.com/en-us/security/business/security-101/what-is-data-loss-prevention-dlp [26.4.2024]
[164] Ibid.
[165] https://learn.microsoft.com/en-us/purview/endpoint-dlp-learn-about [26.4.2024]
[166] https://learn.microsoft.com/en-us/purview/dlp-chrome-learn-about [26.4.2024]
[167] https://learn.microsoft.com/en-us/purview/dlp-learn-about-dlp, https://learn.microsoft.com/en-us/purview/dlp-policy-reference

time. The monitoring focuses on activities that involve sharing information that was declared sensitive or confidential with external parties. Alerts can be displayed both in the DLP system and in Microsoft Defender.[168] Organizations can then use the **activity explorer** to further investigate activity metadata and the **content explorer** to review the contents associated with the alert,[169] both of which are also available in the insider risk system (see section 4.1). The DLP system can not only *detect* sensitive data sharing but also prevent "prohibited activities" from happening in the first place.[170] It can, for example, block sensitive information from being shared in a Teams chat or show a popup in Word or other Office applications that notifies employees that they are "engaging in a risky behavior".[171]

Data governance. In addition, Microsoft offers functionality to screen, classify and govern all information stored across an organization, which includes the detection of personal data and other sensitive information. Purview's "information protection system" promises to help organizations "discover, classify, and protect sensitive information wherever it lives or travels".[172] The information protection system can automatically identify and classify sensitive information based on keywords, patterns or AI-based classifiers across Microsoft 365, Azure, cloud-based data stores and third-party applications,[173] from personal data (e.g. credit card numbers) to confidential information and intellectual property (e.g. financial documents, software source code).[174] Purview's data classification functionality can "label" information with respect to "sensitivity" and "retention".[175] Other data governance tools in Purview include the **data map** and the newer **data catalog**.[176]

4.6 Automated decisions on employees based on behavioral risk scores

The Purview insider risk management system can share data about the current "risk level" of users and employees with other Microsoft systems, which then automatically restrict them from performing certain actions.[177] It can, for example, share information about employees whose risk level was assessed as "elevated" or "moderate" by the insider risk system with Purview's data loss prevention (DLP) system, which can then block those employees from sharing sensitive information with persons outside the organization or from copying files to the clipboard.[178] Optionally, it can also raise a DLP policy alert or show the affected employees warnings or "policy tips and education on best practices of handling sensitive data" when they perform certain activities.[179]

This functionality, which Microsoft refers to as **adaptive protection**, should "help organizations automate their response to insider risks and reduce the time required to identify and remediate potential threats". It utilizes data on "risk levels" calculated from "user insights" based on the "analysis of both content and user activities" and "machine learning" to "enforce effective controls on high-risk users while others maintain productivity". The insider risk

[168] https://learn.microsoft.com/en-us/purview/dlp-alert-investigation-learn
[169] https://learn.microsoft.com/en-us/purview/dlp-alert-investigation-learn [26.4.2024]
[170] https://learn.microsoft.com/en-us/purview/dlp-learn-about-dlp [26.4.2024]
[171] Ibid.
[172] https://learn.microsoft.com/en-us/purview/information-protection [26.4.2024]
[173] https://learn.microsoft.com/en-us/purview/dlp-learn-about-dlp [26.4.2024]
[174] https://learn.microsoft.com/en-us/purview/dlp-get-started-with-the-default-policy [26.4.2024]
[175] https://learn.microsoft.com/en-us/purview/data-classification-overview [26.4.2024]
[176] https://learn.microsoft.com/en-us/purview/governance-solutions-overview [26.4.2024]
[177] https://learn.microsoft.com/en-us/purview/insider-risk-management-adaptive-protection [26.4.2024]
[178] https://learn.microsoft.com/en-us/purview/dlp-adaptive-protection-learn [26.4.2024]
[179] https://learn.microsoft.com/en-us/purview/insider-risk-management-adaptive-protection [26.4.2024]

system can also share data about user and employee risk levels with Microsoft's identity and access management system Entra, which can then, for example, block them from accessing certain applications or any application.[180]

4.7 Data sources and categories analyzed by Microsoft Purview

Microsoft Purview's insider risk, communication compliance and eDiscovery systems can utilize data from systems provided by Microsoft and by other vendors, based on a variety of data integration mechanisms.

The **communication compliance system** can monitor and analyze emails, chats, documents, images, meeting transcripts and other contents in Microsoft 365, Exchange, Teams, Viva Engage and Copilot,[181] including Copilot prompts and responses.[182] In addition, employers can use "data connectors" to include data from third-party communication systems (e.g. Zoom, Slack, Cisco Webex, RingCentral) and messaging software specifically used in the financial sector (e.g. Instant Bloomberg, ICE Chat). In partnership with TeleMessage, the communication compliance system can record and monitor encrypted conversations directly from mobile apps (e.g. WhatsApp, Signal).[183] Employers can let TeleMessage record both voice calls and messages directly from the employees' Android phones or even access call data via partnerships with mobile carrier networks such as AT&T, Verizon, O2 and T-Mobile.[184] Any other data can be imported via custom connectors (e.g. EML mailboxes, SQL databases, CSV, XML).[185]

Microsoft's **eDiscovery** system can access conversations and documents from the same sources as the communication compliance system.[186] This can include emails, chats, documents, images, recorded video meetings and other content in Microsoft 365, Exchange, Teams, SharePoint, Viva Engage and Microsoft Copilot.[187] In addition, it can archive and search data from third-party systems (e.g. Zoom, Slack, Cisco Webex, Bloomberg),[188] and, via third-party integrations, from mobile devices (e.g. voice calls and SMS, WhatsApp messages and calls).[189]

The **insider risk system** can access log data from Microsoft 365 via built-in "audit logs",[190] which represent a "summary of all activities" in an organization,[191] and via the Microsoft Graph API,[192] which the company refers to as the "gateway to data and intelligence in Microsoft 365". Employers can use the Graph API to "access the tremendous amount of data in Microsoft 365, Windows, and Enterprise Mobility + Security".[193] This includes data from Exchange, SharePoint, Teams, OneDrive and other Microsoft 365 applications.[194] The insider risk system can

---

[180] Ibid.

[181] https://learn.microsoft.com/en-us/purview/communication-compliance, https://learn.microsoft.com/en-us/purview/communication-compliance-channels [18.3.2024]

[182] https://learn.microsoft.com/en-us/purview/communication-compliance-copilot?source=recommendations [25.4.2024]

[183] https://learn.microsoft.com/en-us/purview/archive-third-party-data [18.3.2024]

[184] https://learn.microsoft.com/en-us/purview/archive-third-party-data, https://learn.microsoft.com/en-us/purview/archive-enterprise-number-data, https://learn.microsoft.com/en-us/purview/archive-o2-network-data, https://www.telemessage.com/, https://www.telemessage.com/mobile-archiver/network-archiver/, https://www.telemessage.com/category/faq/archiver/network-archiver/ [18.3.2024]

[185] https://learn.microsoft.com/en-us/purview/archive-third-party-data [18.3.2024]

[186] https://learn.microsoft.com/en-us/purview/archive-third-party-data [18.3.2024]

[187] https://learn.microsoft.com/en-us/purview/ediscovery, https://learn.microsoft.com/en-us/purview/ediscovery-overview, https://techcommunity.microsoft.com/t5/security-compliance-and-identity/keep-microsoft-teams-meetings-compliant-with-communication/ba-p/3933446 [25.4.2024]

[188] https://learn.microsoft.com/en-us/purview/archive-third-party-data [25.4.2024]

[189] https://learn.microsoft.com/en-us/purview/archive-enterprise-number-data, https://learn.microsoft.com/en-us/purview/archive-whatsapp-data [25.4.2024]

[190] See also section 4.9

[191] https://learn.microsoft.com/en-us/purview/insider-risk-management-configure [18.3.2024]

[192] https://learn.microsoft.com/en-us/purview/insider-risk-management-solution-overview [18.3.2024]

[193] https://learn.microsoft.com/en-us/graph/overview [18.3.2024]

[194] https://learn.microsoft.com/en-us/purview/insider-risk-management-configure, https://learn.microsoft.com/en-us/purview/insider-risk-management-content-explorer [25.4.2024]

also access data from Microsoft's identity and access management system Entra, formerly known as Azure Active Directory,[195] including information about an employee's roles, permissions, job title and position in the organizational hierarchy.[196] It can also import "employee profile data" from HR systems.[197] Furthermore, it can utilize activity data from the computers and devices used by employees.[198] Windows devices can be monitored either via Microsoft's antivirus solution "Defender for Endpoint",[199] the company's mobile device management (MDM) tool Intune[200] and other "onboarding" technologies.[201] Microsoft also provides ways to access data from macOS devices.[202] To access data on web browsing, the "Microsoft Compliance Extension" for the Edge or Chrome web browser can be installed on employees' devices.[203]

Based on data from Microsoft software, the insider risk system can monitor, for example, login, file, printing, clipboard, application, browser, meeting, email and chat activities.[204] To monitor the contents of communication activities, it can incorporate data from the communication compliance system.[205] When the "forensic evidence" functionality is activated, it can record screen activity and every user interaction performed on employee devices.[206] In addition to data about user account and job title changes from Microsoft Entra, the system can incorporate performance reviews, performance improvement plans, terminations and other human resource information via the "HR data connector".[207] It can, for example, access data on a "below average" performance rating including the date the corresponding employee was informed about the result of their performance review.[208] To monitor physical movements in buildings, offices and other facilities, it can access badging data.[209]

More broadly, the insider risk system can incorporate activity data from almost any other enterprise software system. Organizations can either utilize "data connectors" provided by Microsoft and other vendors[210] or import "risk indicator" and alert data from SIEM systems such as Microsoft Sentinel or Splunk, which continuously aggregate activity log data from various enterprise software systems. As such, this can include any activity and alert data processed by Sentinel (e.g. network and firewall activity, see section 4.8.1). Microsoft's software documentation describes how its insider risk system can monitor suspicious reporting and file activities performed in Salesforce and Dropbox.[211] In addition, the system can import activity data from Microsoft's "Defender for Cloud Apps" system,[212] which monitors more than 31,000 different cloud-based applications for cybersecurity purposes.[213]

---

[195] https://learn.microsoft.com/en-us/entra/identity/ [18.3.2024]
[196] https://learn.microsoft.com/en-us/purview/insider-risk-management-settings-policy-indicators [18.3.2024]
[197] https://learn.microsoft.com/en-us/purview/import-hr-data [26.3.2024]
[198] https://learn.microsoft.com/en-us/purview/insider-risk-management-settings-policy-indicators [18.3.2024]
[199] Ibid.
[200] https://learn.microsoft.com/en-us/purview/device-onboarding-mdm [18.3.2024]
[201] https://learn.microsoft.com/en-us/purview/device-onboarding-overview [18.3.2024]
[202] https://learn.microsoft.com/en-us/purview/device-onboarding-macos-overview [18.3.2024]
[203] https://learn.microsoft.com/en-us/purview/insider-risk-management-browser-support [7.3.2024]
[204] See section 4.1
[205] See section 4.3
[206] See section 4.1.5
[207] https://learn.microsoft.com/en-us/purview/insider-risk-management-configure [18.3.2024]
[208] https://learn.microsoft.com/en-us/purview/import-hr-data [26.3.2024]
[209] https://learn.microsoft.com/en-us/purview/import-physical-badging-data [18.3.2024]
[210] https://learn.microsoft.com/en-us/purview/archive-third-party-data [18.3.2024]
[211] https://learn.microsoft.com/en-us/purview/import-insider-risk-indicators [18.3.2024]
[212] https://learn.microsoft.com/en-us/purview/purview-data-flows [18.3.2024]
[213] https://learn.microsoft.com/en-us/defender-cloud-apps/editions-cloud-app-security-o365 [18.3.2024]

4.8 Analyzing activity log data for cybersecurity and other purposes with Microsoft Sentinel

While Microsoft Purview focuses on risk and compliance, the company's "Sentinel" system addresses cybersecurity at large. **Microsoft Sentinel** combines "security information and event management" (SIEM) with "user and entity behavior analytics" (UEBA) and also provides "security orchestration, automation and response" (SOAR) functionality. It promises to help organizations detect cybersecurity threats and investigate them with the help of "artificial intelligence".[214] At its core, it collects and analyzes large amounts of log data from many different sources "across all users, devices, applications, and infrastructure, both on-premises and in multiple clouds".[215]

Figure 15 (top left) shows Sentinel's "overview" dashboard. In this example, the system has collected and analyzed 7,400 activity records from different data sources, which are referred to as "events". The data sources include log files that record activity in the organization's Microsoft 365 and Azure[216] systems, including sign-in activity.

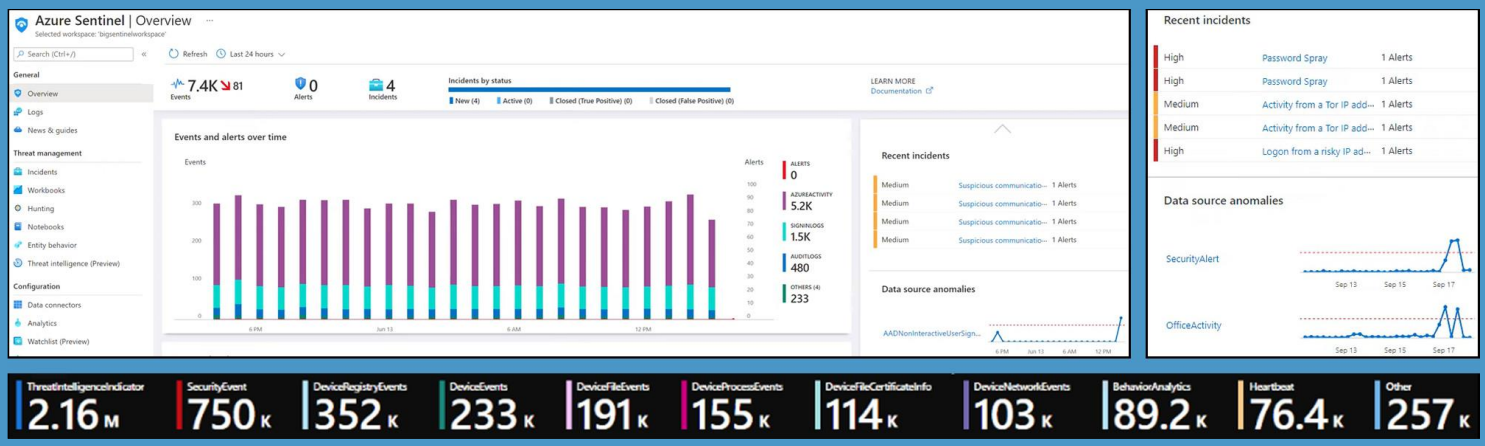

Figure 15: Monitored activities, event log data sources, alerts, incidents and anomalies (Microsoft Sentinel)[217]

The dashboard also displays information about alerts and recent incidents, for example, about alerts that involve user passwords and activities related to "risky" IP addresses (Figure 15, top right) or "suspicious communication" activity (Figure 15, top left). Furthermore, it displays information about "data source anomalies" indicating, for example, a recent rise in the number of "office activity" events (Figure 15, top right).

Large amounts of personal data about employee behavior. The number of event log records processed by Sentinel can be very large. Figure 15 (bottom) shows a part of the user interface of another example Sentinel system that processed several million events from different activity logs. This includes, for example, activity logs named "device file events", "device process events" and "device network events", which provide large amounts of data about how employees use files, applications and network connectivity on their devices. The data is collected via

[214] https://learn.microsoft.com/en-us/azure/sentinel/overview, https://learn.microsoft.com/en-us/azure/sentinel/identify-threats-with-entity-behavior-analytics [20.3.2024]
[215] Ibid.
[216] "Azure" refers to Microsoft's cloud platform at large, https://azure.microsoft.com [18.3.2024]
[217] Figures © Microsoft, T-Minus 365. The figures serve as basis for the discussion of the corporate practices examined in this study. Sources: video min 0:01: https://www.youtube.com/watch?v=EIywOXnfmkk, video min 0:20: https://www.youtube.com/watch?v=3JLsJPsy8m8, https://techcommunity.microsoft.com/t5/microsoft-sentinel-blog/run-microsoft-sentinel-playbooks-from-workbooks-on-demand/ba-p/3193074 [30.10.2023]

Microsoft's antivirus software, "Defender for Endpoint", installed on the employees' devices.[218] The "device file events" log, for example, contains activity records about the creation and modification of files. It includes information about the concerned file, the time the activity was performed and the device and user who initiated it.[219] Similarly, the "device process events" log contains activity records about application usage, which include information about the device and user who initiated the activity.[220] As such, the Sentinel system processes large amounts of personal data about employee behavior.

These example dashboards also demonstrate that Sentinel can collect and analyze **very different types of activity records** on employee behavior ranging from data related to sign-ins, passwords and IP addresses to records on device, file, application, network, communication and office activity. Figure 16 (left) shows a Sentinel report about **Office 365 activity** over a period of time that addresses 1,040 activities in Microsoft Exchange, 274 activities in SharePoint, 60 activities in OneDrive and 6 activities in Microsoft Teams.

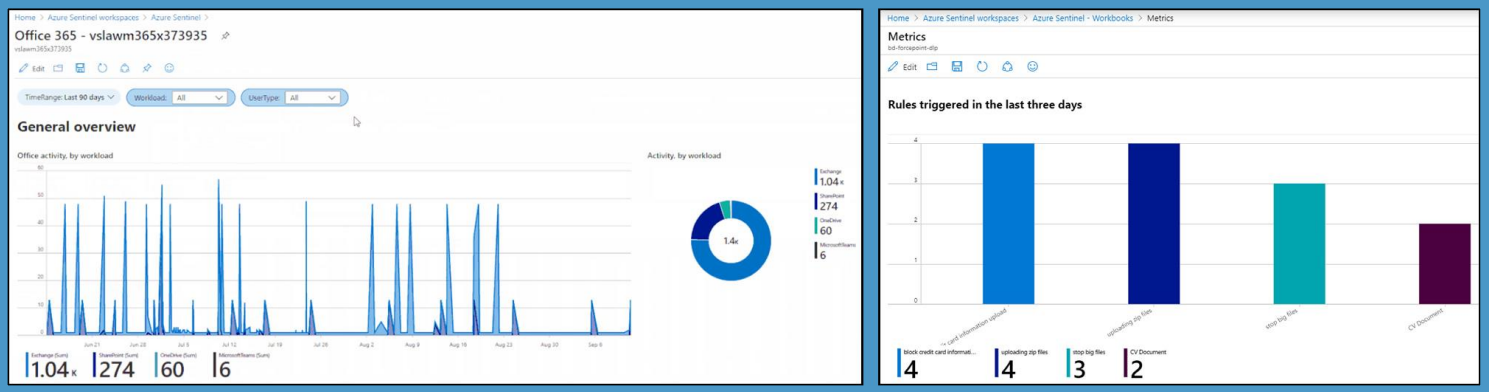

Figure 16: Office 365, Exchange, SharePoint, OneDrive, Teams and Forcepoint activity data (Microsoft Sentinel)[221]

Sentinel monitors activities associated with different **entities** such as user accounts, devices (hosts), applications (processes), files, IP addresses, URLs, Azure resources, mailboxes and mail messages. User accounts, devices and other entities often contain identifying information that refers to employees.[222]

4.8.1 Sentinel data sources

Microsoft offers hundreds of "data connectors"[223] that allow organizations to import data from different enterprise systems from different vendors into Sentinel and then continuously monitor activities.[224] It provides 48 data connectors to access data from **Microsoft systems** such as Azure, Entra, Purview, Microsoft 365, Dynamics 365, Project, GitHub, Azure and Windows Firewall and Defender, including Defender for Endpoint, Identity, IoT and

[218] https://learn.microsoft.com/en-us/azure/sentinel/connect-microsoft-365-defender?tabs=MDE [20.3.2024]
[219] https://learn.microsoft.com/en-us/microsoft-365/security/defender/advanced-hunting-devicefileevents-table [20.3.2024]
[220] https://learn.microsoft.com/en-us/microsoft-365/security/defender/advanced-hunting-deviceprocessevents-table [20.3.2024]
[221] Figures © Microsoft, Forcepoint. The figures serve as basis for the discussion of the corporate practices examined in this study. Sources: https://techcommunity.microsoft.com/t5/microsoft-sentinel-blog/how-to-protect-office-365-with-azure-sentinel/ba-p/1656939, https://forcepoint.github.io/docs/dlp_and_azure_sentinel/ [30.10.2023]
[222] https://learn.microsoft.com/en-us/azure/sentinel/entities-reference [28.3.2024]
[223] https://learn.microsoft.com/en-us/azure/sentinel/data-connectors-reference [20.3.2024]
[224] https://learn.microsoft.com/en-us/azure/sentinel/connect-data-sources [20.3.2024]

Cloud.[225] Microsoft explains that Sentinel can access "most telemetry and event data".[226] Purview's insider risk and communication compliance systems, as examined in sections 4.1 and 4.2, can also share alert data about suspicious activities with Sentinel.[227]

In addition, a number of data connectors to **systems from third-party vendors** are available, for example:[228]

- Cloud-based enterprise software systems and environments, e.g. Salesforce, Oracle, SAP, Atlassian Confluence/Jira, Zoom, Slack, Google Workspace, Workplace for Facebook, Snowflake
- Large-scale cloud environments, e.g. Amazon Web Services, Google Cloud, AliCloud, Oracle Cloud
- Cybersecurity systems and IT infrastructure including networking technology, virtualization, device/antivirus protection, SIEM, UEBA, threat intelligence and CDN, e.g. Cisco, Juniper, Okta, Citrix, VMware, Symantec, McAfee, Kaspersky, Forcepoint, Exabeam, Fortinet, Rapid7, Proofpoint, SentinelOne, Palo Alto Networks, Digital Guardian, Greynoise Intelligence, Crowdstrike, Akamai, Cloudflare
- Other technologies, e.g. Linux Syslog, RedHat JBoss, Apache, Nginx, PostgreSQL, MongoDB

All these systems can share either data about all activities with Microsoft Sentinel or only some data, for example, alerts about suspicious activities. Organizations can, for example, import Salesforce event logs[229], Cisco access logs[230], Zoom reports[231], Oracle Cloud event logs[232] or Jira "audit" logs[233] into their Sentinel system. Microsoft provides extra functionality to import extensive log data from the major ERP system SAP.[234] Via custom "data connectors", Sentinel can access any other data.[235]

Organizations that use cybersecurity solutions from **Forcepoint**, as examined in section 3, can import data from Forcepoint's firewall,[236] cloud security,[237] access security[238] and data loss prevention (DLP)[239] systems.[240] According to Microsoft, importing Forcepoint data "enriches visibility into user activities across locations and cloud applications, enables further correlation with data from Azure […] and other feeds, and improves monitoring capability […] inside Microsoft Sentinel".[241] Figure 16 (right) demonstrates how Sentinel displays information about suspicious activities imported from Forcepoint's risk profiling system. In the previous three days, it detected the upload of "zip files" and the presence of "CV documents", both of which were assessed as suspicious activities by Forcepoint. The imported incident records can now be used for further analysis in Microsoft Sentinel.

[225] https://learn.microsoft.com/en-us/azure/sentinel/data-connectors-reference [20.3.2024]
[226] https://learn.microsoft.com/en-us/azure/sentinel/hunting [20.3.2024]
[227] https://learn.microsoft.com/en-us/purview/insider-risk-management-settings-dlp-sync, https://learn.microsoft.com/en-us/purview/communication-compliance-siem [20.3.2024]
[228] https://learn.microsoft.com/en-us/azure/sentinel/data-connectors-reference [20.3.2024]
[229] https://learn.microsoft.com/en-us/azure/sentinel/data-connectors/salesforce-service-cloud-using-azure-functions [20.3.2024]
[230] https://learn.microsoft.com/en-us/azure/sentinel/data-connectors/cisco-web-security-appliance [20.3.2024]
[231] https://learn.microsoft.com/en-us/azure/sentinel/data-connectors/zoom-reports-using-azure-functions [20.3.2024]
[232] https://learn.microsoft.com/en-us/azure/sentinel/data-connectors/oracle-cloud-infrastructure-using-azure-functions [20.3.2024]
[233] https://learn.microsoft.com/en-us/azure/sentinel/data-connectors/atlassian-jira-audit-using-azure-functions [20.3.2024]
[234] https://learn.microsoft.com/en-us/azure/sentinel/sap/solution-overview [20.3.2024]
[235] https://learn.microsoft.com/en-us/azure/sentinel/create-custom-connector [20.3.2024]
[236] https://learn.microsoft.com/en-us/azure/sentinel/data-connectors/recommended-forcepoint-ngfw-via-ama [20.3.2024]
[237] https://learn.microsoft.com/en-us/azure/sentinel/data-connectors/recommended-forcepoint-csg-via-ama [20.3.2024]
[238] https://learn.microsoft.com/en-us/azure/sentinel/data-connectors/recommended-forcepoint-casb-via-ama [20.3.2024]
[239] https://learn.microsoft.com/en-us/azure/sentinel/data-connectors/forcepoint-dlp [20.3.2024]
[240] https://learn.microsoft.com/en-us/azure/sentinel/data-connectors-reference [20.3.2024]
[241] https://learn.microsoft.com/en-us/azure/sentinel/data-connectors/recommended-forcepoint-casb-via-ama [20.3.2024]

4.8.2 Profiling, ranking and singling out employees with Microsoft's UEBA technology

Sentinel offers functionality for "user and entity behavior analytics" (UEBA), which promises to identify "anomalous" behavior that may indicate a cybersecurity threat.[242] It can be used to "capture non-routine actions" and find "potentially non-compliant practices", according to Microsoft.[243] Based on "logs and alerts" from Sentinel's data sources, the UEBA system creates "baseline behavioral profiles" for employees and other "entities" (e.g. hosts, devices, files, applications) and then aims to detect behavior that deviates from these profiles. Microsoft emphasizes that its UEBA technology aims to identify both "compromised entities" and "malicious insiders".[244] Similar to Microsoft's insider risk system, it addresses both hacked user accounts and employees who are considered risks.

Ranking employees by risk. Figure 17 (top left) shows how the system displays a list of "top users to investigate", ranked by the number of associated incidents, alerts and detected "anomalies". It lists named employees who are considered risks, including their email addresses. Figure 17 (bottom left) shows a list of anomalies detected for a specific employee who performed the activities "anomalous login to device" and "anomalous resource access", the latter of which refers to unusual access to resources such as storage accounts, databases or applications.[245]

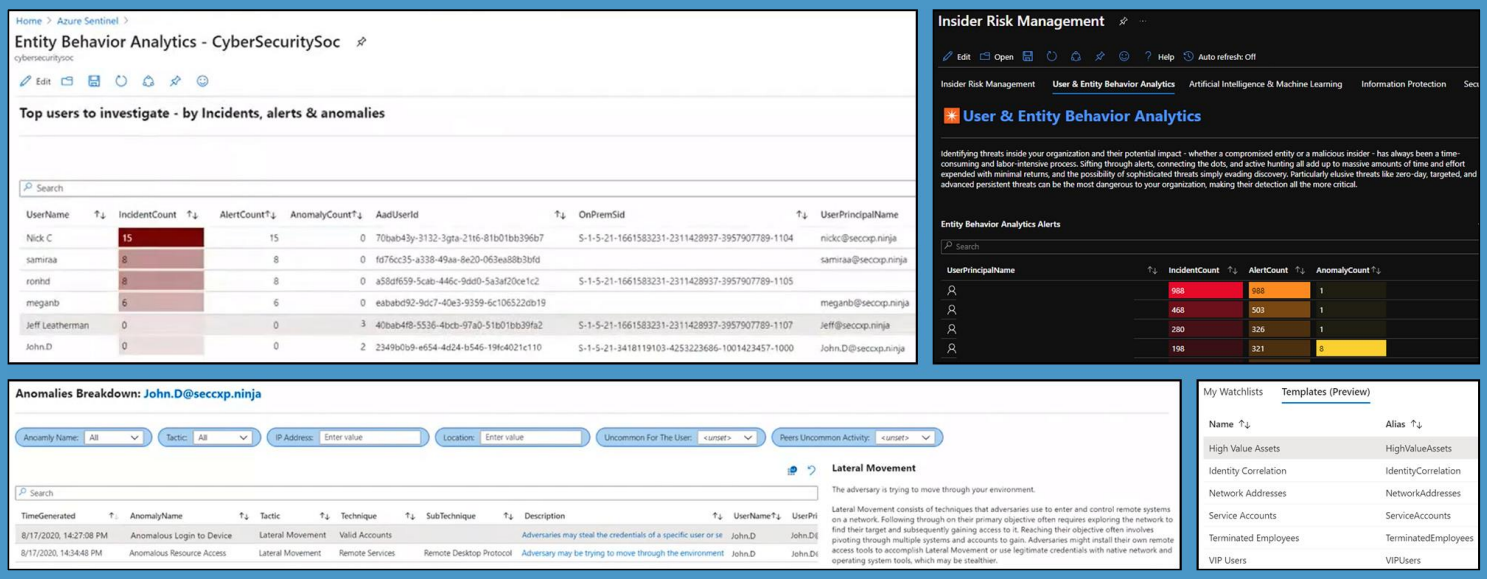

Figure 17: Ranking employees by risk, suspicious activities and watchlists (Microsoft Sentinel/Purview) [246]

In this example, the system recognized these activities as potential parts of a "lateral movement" threat, which refers to a cyberattack where an attacker slowly moves through an organization's IT system. While these activities appear

[242] https://learn.microsoft.com/en-us/azure/sentinel/identify-threats-with-entity-behavior-analytics [25.3.2024]

[243] https://learn.microsoft.com/en-us/azure/sentinel/investigate-with-ueba [25.3.2024], https://techcommunity.microsoft.com/t5/microsoft-sentinel-blog/guided-ueba-investigation-scenarios-to-empower-your-soc/ba-p/1857100 [25.3.2024]

[244] https://learn.microsoft.com/en-us/azure/sentinel/identify-threats-with-entity-behavior-analytics [25.3.2024]

[245] https://learn.microsoft.com/en-us/entra/architecture/secure-resource-management [25.3.2024]

[246] Figures © Microsoft. The figures serve as basis for the discussion of the corporate practices examined in this study. Sources: video min 32:56 and 35:56: https://www.youtube.com/watch?v=IWGTc-yQ9FY, https://techcommunity.microsoft.com/t5/microsoft-sentinel-blog/announcing-the-microsoft-purview-insider-risk-management/ba-p/2955786, https://learn.microsoft.com/en-us/azure/sentinel/watchlists-create [30.10.2023]

to address external cyberattacks, Sentinel's UEBA functionality can also be utilized in Microsoft's insider risk system.[247] Figure 17 (top right) shows how a list of suspicious employees ranked by the number of alerts and anomalies is displayed in the insider risk system. Organizations can potentially use the UEBA system to target employees.

Identifying "anomalous" behavior. The UEBA system utilizes data about users, devices and activities from sources such as Microsoft Entra, Azure activity logs and Windows security events.[248] It can "correlate security events" with identity records[249] by combining activity log records with Entra information such as employee name, email address, phone number, job title, department, assigned roles and group memberships.[250] To detect anomalous behavior, the system constantly monitors activities over a period of time to establish a "baseline of legitimate activity" and then considers "any activity outside the normal parameters" as "anomalous and therefore suspicious".[251] The baseline profiles are calculated based on "behavioral learning" [252] about activities performed by employees, their "peers" and those of the "organization as a whole".[253] When the system detects anomalies, it assigns each suspicious activity an "investigation priority score". Activities that are "identified as the most abnormal" receive the highest scores.[254] To identify an employee's "peers", Sentinel uses Entra information to create a ranked list of up to 20 other users.[255] Sentinel utilizes various machine learning models to detect different types of anomalies. In addition to anomaly detection based on UEBA profiling, it offers "customizable" anomaly detection templates. The following list provides examples of anomalous activities that Sentinel promises to detect out of the box:[256]

- Anomalous sign-in
- Anomalous account creation
- Login from an unusual region
- Anomalous user activities in Office Exchange
- Anomalous user/app activities in Azure audit logs
- Anomalous web request activity
- Unusual network volume anomaly
- Excessive data transfer anomaly
- Anomalous data destruction
- Suspicious number of protected documents accessed

The UEBA system potentially creates **large numbers of records about behavioral anomalies**. Figure 15 (bottom) shows how Sentinel processes 89,200 "behavior analytics" records. The report in Figure 15 (top right) displays a chart that indicates the level of anomalous activities in Office. Sentinel's anomaly detection templates "were developed to be robust by using thousands of data sources and millions of events". Microsoft emphasizes that anomalies are "notoriously very noisy" and "typically require a lot of tedious tuning for specific environments". While a "single anomaly is not a strong signal of malicious behavior", several anomalies that occur at different points may

[247] https://techcommunity.microsoft.com/t5/microsoft-sentinel-blog/announcing-the-microsoft-purview-insider-risk-management/ba-p/2955786 [25.3.2024]
[248] https://learn.microsoft.com/en-us/azure/sentinel/ueba-reference [25.3.2024]
[249] https://learn.microsoft.com/en-us/azure/sentinel/investigate-with-ueba [25.3.2024]
[250] https://learn.microsoft.com/en-us/azure/sentinel/ueba-reference [25.3.2024]
[251] https://learn.microsoft.com/en-us/azure/sentinel/anomalies-reference [25.3.2024]
[252] https://learn.microsoft.com/en-us/azure/sentinel/identify-threats-with-entity-behavior-analytics [25.3.2024]
[253] https://learn.microsoft.com/en-us/azure/sentinel/anomalies-reference [25.3.2024]
[254] https://learn.microsoft.com/en-us/azure/sentinel/identify-threats-with-entity-behavior-analytics [25.3.2024]
[255] https://learn.microsoft.com/en-us/azure/sentinel/identify-threats-with-entity-behavior-analytics [25.3.2024]
[256] https://learn.microsoft.com/en-us/azure/sentinel/anomalies-reference [25.3.2024]

cumulatively raise an alert in the system.[257] To detect "advanced multistage attacks", Sentinel can trigger "incidents" by correlating "multiple signals" from various data sources. For example, it can trigger an "data exfiltration" incident by detecting a "suspicious sign-in" followed by the download of an "anomalous number of files" or "sharing files such as documents, spreadsheets, etc., with unauthorized users".[258]

4.8.3 Putting employees on "watchlists" and investigating past activity and relationships

To focus on certain types of entities or employees, Sentinel allows for the creation of "watchlists".[259] As shown in Figure 17 (bottom right), Microsoft provides a number of pre-configured watchlists. The "high value assets" watchlist, for example, addresses activities that involve certain "devices, resources, and other assets that have critical value in the organization". The "network addresses" watchlist focuses on activities that involve certain IP addresses and their "organizational contexts". The "VIP users" watchlist addresses "user accounts of employees that have high impact value in the organization". The "terminated employees" watchlist focuses on activities that involve "user accounts of employees that have been, or are about to be, terminated".[260] Employers can create custom watchlists based on certain search queries.[261]

Investigating past employee activity and relationships with other entities. As demonstrated in Figure 18 (left), Sentinel can display information about suspicious activities for a certain named employee over a period of time including alerts, anomalies and other "events" such as Office activities. The system can also list "similar" incidents and related entities,[262] such as other employees, accounts, hosts, IP addresses, processes, files and mailboxes.[263] The "investigation graph" provides a visual representation of relationships between entities.[264]

Figure 18: Investigating past employee activity and entity relationships in detail (Microsoft Sentinel)[265]

[257] https://learn.microsoft.com/en-us/azure/sentinel/work-with-anomaly-rules [25.3.2024]
[258] https://learn.microsoft.com/en-us/azure/sentinel/fusion-scenario-reference [25.3.2024]
[259] https://learn.microsoft.com/en-us/azure/sentinel/watchlists [25.3.2024]
[260] https://learn.microsoft.com/en-us/azure/sentinel/fusion-scenario-reference [25.3.2024]
[261] https://learn.microsoft.com/en-us/azure/sentinel/watchlists-queries [25.3.2024]
[262] https://learn.microsoft.com/en-us/azure/sentinel/investigate-incidents [25.3.2024]
[263] https://learn.microsoft.com/en-us/azure/sentinel/entities [25.3.2024]
[264] https://learn.microsoft.com/en-us/azure/sentinel/investigate-incidents [25.3.2024]
[265] Figures © Microsoft, Thomas Naunheim. The figures serve as basis for the discussion of the corporate practices examined in this study. Sources: video min 15:49: https://www.youtube.com/watch?v=j8mhviNFfSI, https://www.cloud-architekt.net/identity-security-monitoring/ [30.10.2023]

Figure 18 (right) shows an example investigation graph that displays relationships between an employee user account, an IP address, OneDrive and a number of suspicious activities such as failed logins and file accesses.

4.8.4 Ranking employees by risk and investigating past activity – UEBA in Microsoft Defender

Microsoft's UEBA technology[266] can not only be utilized in its Sentinel system but also in "Microsoft 365 Cloud App Security", which is part of the company's "Defender for Cloud Apps" product.[267]

Figure 19 (bottom center) shows how Defender displays a list of suspicious employees ranked by their "investigation priority score". The report in Figure 19 (left) displays detailed information about one of these suspicious employees, who was assigned a risk score of 155. According to the report, 13 alerts and other "risky activities" contributed to the calculation of the score. This includes activity from a Tor[268] IP address, the download of a certain file and access to a specific file server. The employee's risk score is compared to other employees across the organization. Subsequently, investigators can "dive deeper into each one of the alerts or activities".[269]

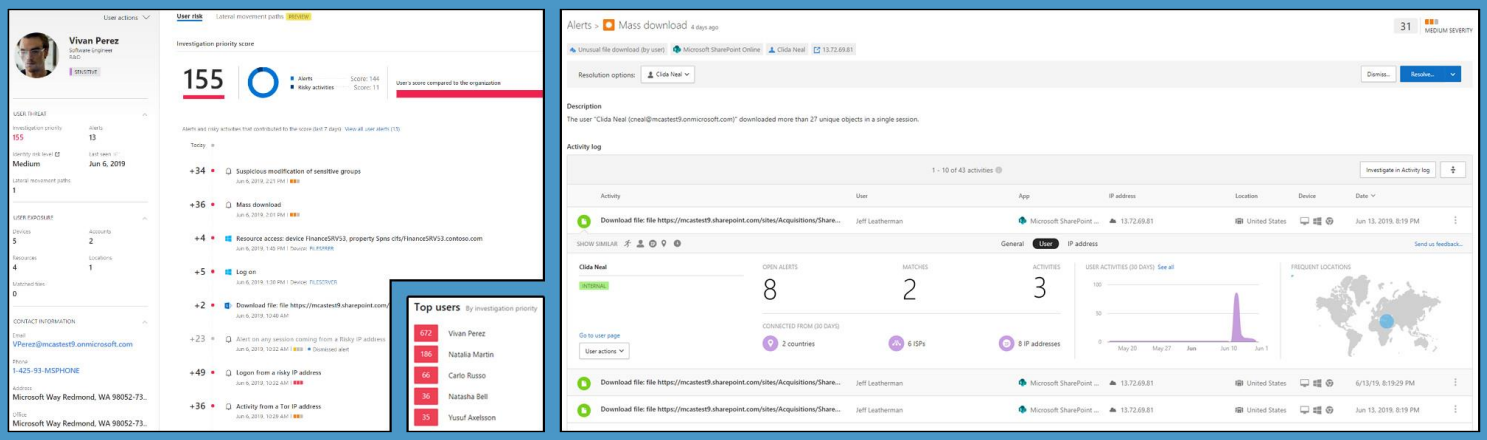

Figure 19: Singling out employees, ranking them by risk and investigating past activities (Microsoft Defender)[270]

Figure 19 (right) shows another report about a named employee who downloaded 27 documents from an organization's SharePoint system. The list contains detailed information about each download activity, including the file name and the date and time the download was performed. The user has "8 open alerts" and connected to the system from two countries and eight IP addresses in the last 30 days. While this report appears to address an external cyberattack involving user accounts of employees, the same functionality can potentially be used to investigate employee activities for other purposes. The example report in Figure 20 (left) shows another list of named employees ranked by their risk scores as determined by Microsoft's UEBA technology in the company's "Defender for Cloud Apps" product.

---

[266] The software documentation for Sentinel's UEBA technology refers to the use of the same technology in Microsoft Defender, e.g. "See how behavior analytics is used in Microsoft Defender for Cloud Apps for an example of how this works", https://learn.microsoft.com/en-us/azure/sentinel/identify-threats-with-entity-behavior-analytics [28.3.2024]

[267] https://learn.microsoft.com/en-us/defender-cloud-apps/editions-cloud-app-security-o365 [28.3.2024]

[268] Tor is a tool for anonymous web surfing used by journalists, whistleblowers and cybercriminals: https://en.wikipedia.org/wiki/Tor_(network)

[269] https://techcommunity.microsoft.com/t5/security-compliance-and-identity/prioritize-user-investigations-in-cloud-app-security/ba-p/700136 [28.3.2024]

[270] Figures © Microsoft. The figures serve as basis for the discussion of the corporate practices examined in this study. Source: https://techcommunity.microsoft.com/t5/security-compliance-and-identity/prioritize-user-investigations-in-cloud-app-security/ba-p/700136 [30.10.2023]

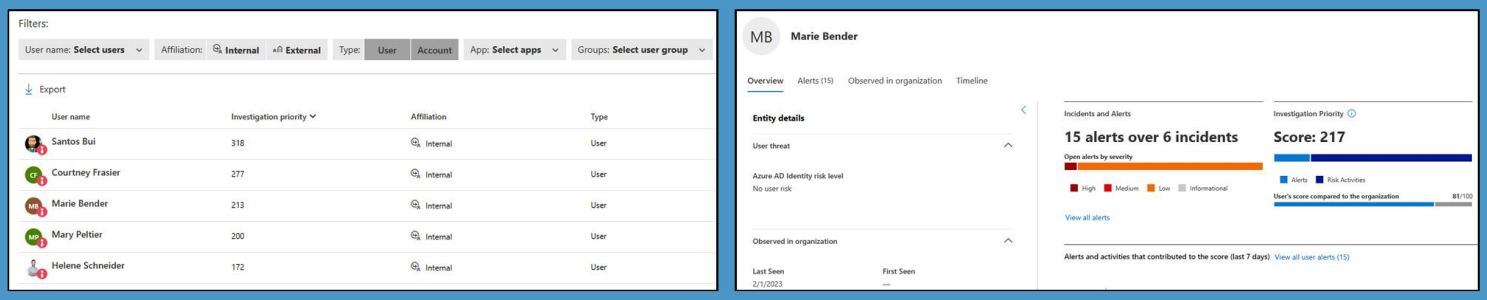

Figure 20: Singling out employees and ranking them by risk (Microsoft Defender)[271]

Figure 20 (right), once again, demonstrates how detailed information about a specific employee is displayed, including about alerts and "risk activities" that contributed to the employee's risk score. According to Microsoft, its UEBA system "builds user profiles for each user based on analytics that take time, peer groups, and expected user activity into consideration".[272] It lets organizations "immediately understand who the real top risky users are" through an evaluation of "anomalous" activity, which is based on "dynamic peer calculations and machine learning" and determines the "investigation priority for each user", according to Microsoft. The system "gives [a score] to each user to let [organizations] know how risky a user is relative to other users in [their] organization". The risk score represents the "sum of all the user's risky activities over the last week".[273]

Identifying "disgruntled employees". Once again, Microsoft's UEBA system promises to detect both "external attackers" and "malicious insiders". After suspicious users or employees are detected, organizations can "investigate all related activities – whether they're compromised, exfiltrating data, or acting as insider threats". They can "get information about who the user is and what is known about them". Microsoft promises to help organizations decide whether a user that was assessed to be a risk is an "engineer who often performs unusual activities as part of their job" or a "disgruntled employee who just got passed over for a promotion".[274]

4.8.5 Custom functionality, apps, queries and "dragnet" searches for employee activity

As a SIEM system, Microsoft Sentinel offers various ways to collect and analyze activity data from many sources.[275] The previous sections describe how it helps organizations single out "anomalous" or otherwise suspicious behaviors, user accounts and employees, rank them by risk and further investigate past activities. Most of the reports and analysis functions examined in sections 4.8.1 to 4.8.4 combine different **Sentinel components** such as "data connectors", "analytics rules", "hunting queries", "workbooks" and "playbooks".[276]

A Sentinel **workbook** provides monitoring and reporting functionality for a specific use case.[277] The "User and Entity Behavior Analytics" workbook,[278] for example, provides UEBA functionality including reports about suspi-

[271] Figures © Microsoft. The figures serve as basis for the discussion of the corporate practices examined in this study. Source: https://learn.microsoft.com/en-us/defender-cloud-apps/tutorial-ueba [30.10.2023]
[272] https://learn.microsoft.com/en-us/defender-cloud-apps/tutorial-ueba [28.3.2024]
[273] Ibid.
[274] Ibid.
[275] See section 4.8.1
[276] https://learn.microsoft.com/en-us/azure/sentinel/sentinel-solutions [15.4.2024]
[277] Ibid.
[278] https://learn.microsoft.com/en-us/azure/sentinel/identify-threats-with-entity-behavior-analytics [15.4.2024]

cious users and employees, as detailed in section 4.8.2. The "Office 365" workbook provides insights into SharePoint, OneDrive, Teams and Exchange activity,[279] as shown in Figure 16. The "insider risk management" workbook contains Sentinel functionality related to the detection and investigation of employees who are considered "insider threats".[280] Workbooks may rely on other Sentinel components such as **data connectors** (to import log data from different enterprise systems)[281] and **analytics rules** (to analyze the data, detect threats and raise alerts).[282] "Automation rules" and **playbooks**, which are both part of Sentinel's "security orchestration, automation and response" (SOAR) functionality, help to automate how to respond to detected threats, for example, by creating lists of tasks or triggering complex workflows across several systems in response to specific alerts and incidents.[283]

Microsoft offers a number of pre-built workbooks[284] and sets of components[285] for different purposes ranging from IT operations to cybersecurity to compliance.[286] Employers can also install Sentinel components provided by third-party vendors via Microsoft's cloud-based **app store**,[287] which the company refers to as the "Azure Marketplace".[288] This includes, for example, Sentinel integrations for enterprise software systems from companies such as Amazon, Google, Oracle, SAP, Salesforce, Atlassian, Zoom, VMware, Blackberry, Cisco, Juniper, McAfee, Symantec and Forcepoint.[289] Sentinel components are customizable.[290] Employers can use either use pre-built components provided by Microsoft or third-party vendors and then customize them or create custom components from scratch.

[279] https://learn.microsoft.com/en-us/azure/sentinel/top-workbooks [15.4.2024]
[280] https://techcommunity.microsoft.com/t5/microsoft-sentinel-blog/announcing-the-microsoft-purview-insider-risk-management/ba-p/2955786 [15.4.2024]
[281] See section 4.8.1
[282] https://learn.microsoft.com/en-us/azure/sentinel/detect-threats-built-in [15.4.2024]
[283] https://learn.microsoft.com/en-us/azure/sentinel/automation [15.4.2024]
[284] https://learn.microsoft.com/en-us/azure/sentinel/top-workbooks [15.4.2024]
[285] https://learn.microsoft.com/en-us/azure/sentinel/sentinel-solutions-catalog [15.4.2024]
[286] https://learn.microsoft.com/en-us/azure/sentinel/sentinel-solutions [15.4.2024]
[287] Ibid.
[288] https://azuremarketplace.microsoft.com/ [15.4.2024]
[289] https://azuremarketplace.microsoft.com/en-us/marketplace/apps?filters=solution-templates&page=1&search=sentinel [15.4.2024]
[290] https://learn.microsoft.com/en-us/azure/sentinel/ci-cd-custom-content [15.4.2024]

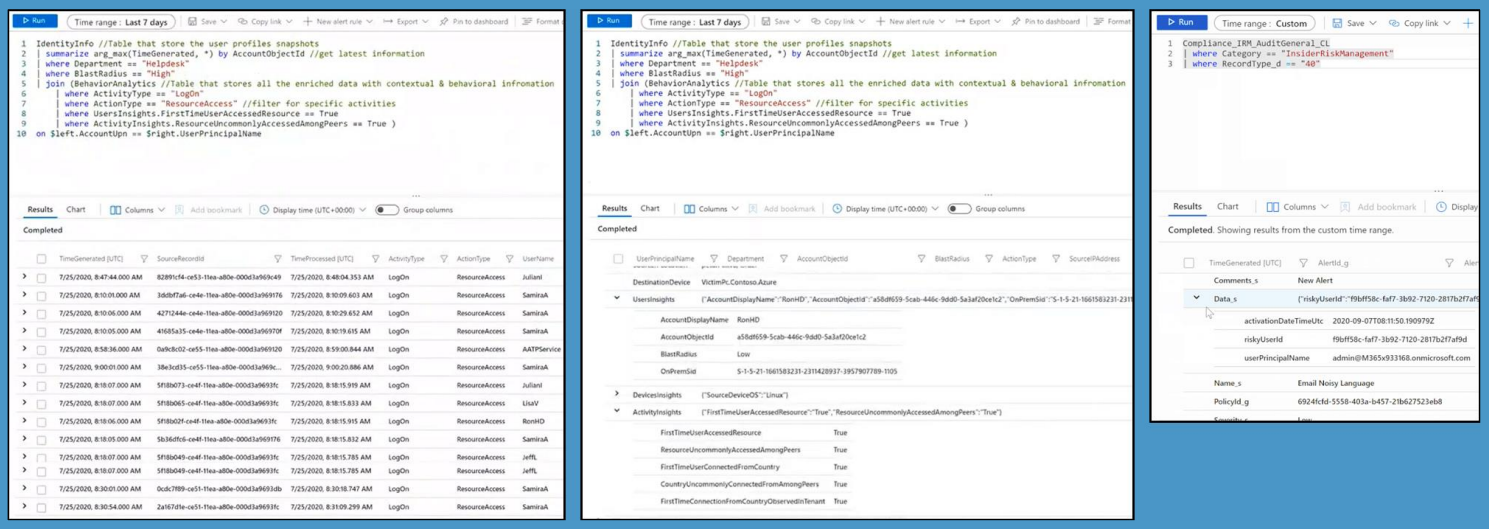

Figure 21: Queries and "dragnet" searches for employee activity (Microsoft Sentinel)[291]

In addition to automated threat detection, alerts and reports, Sentinel offers search and query tools that allow an organization's security analysts and investigators to proactively "hunt" for threats across the available data sources by creating custom **hunting queries** that search for suspicious activities, users and employees.[292] Log and activity data imported to Sentinel is organized into a hierarchy of databases, tables and columns. Microsoft provides an extra programming language similar to the standard database manipulation language SQL, which makes it possible to filter and search the data and which is referred to as the **"Kusto Query Language" (KQL)**.[293] The example KQL query in Figure 21 (left) searches a particular Sentinel database table for users and employees in the "Helpdesk" department who accessed certain "resources" for the first time which are usually not accessed among their "peers". The "results" table displays a list of suspicious resource access activities including the names of users and employees who performed these activities. Figure 21 (center) demonstrates how the system displays details about one of the results including information about the corresponding activity, user and device.

Combining Sentinel, insider risk management and communication monitoring. Figure 21 (right) shows a KQL query that searches a Sentinel database table related to insider risk management for alerts created by Microsoft Purview's "communication compliance" system. The result indicates that the system raised an "email noisy language" alert for a particular employee with a particular username, which refers to the detection of "offensive language" in an email, according to the software documentation.[294] This shows that Sentinel can be used to search for activities that were assessed as suspicious by Microsoft Purview, and thus, for employees who performed certain communication activities. Custom KQL queries can be used to search any data source that is available in Sentinel.[295]

[291] Figures © Microsoft. The figures serve as basis for the discussion of the corporate practices examined in this study. Sources: video min 45:38 and 47:07: https://www.youtube.com/watch?v=IWGTc-yQ9FY, https://techcommunity.microsoft.com/t5/microsoft-sentinel-blog/aggregating-insider-risk-management-information-via-azure/ba-p/1743211 [30.10.2023]

[292] https://learn.microsoft.com/en-us/azure/sentinel/hunting [15.4.2024]

[293] https://learn.microsoft.com/en-us/azure/sentinel/kusto-overview [15.4.2024]

[294] https://techcommunity.microsoft.com/t5/microsoft-sentinel-blog/aggregating-insider-risk-management-information-via-azure/ba-p/1743211 [15.4.2024]

[295] https://learn.microsoft.com/en-us/azure/sentinel/hunting [15.4.2024]

Sentinel makes it possible to search log and activity data from different sources by "joining" databases tables.[296] As such, the system enables extensive **"dragnet" searches for employee activity** according to various criteria.

For "more complex" investigations, organizations can use so-called **notebooks** which provide enhanced data analysis, machine learning and visualization functionality. Microsoft specifically suggests using notebooks to analyze data from external services such as "geolocation data or threat intelligence sources" and "sensitive data" such as "human resource databases or lists of high-value assets".[297] While **machine learning** is "one of the major underpinnings" of Sentinel, according to Microsoft, organizations can also "bring" their "own" machine learning into Sentinel in order to train custom models and then use them to detect suspicious activity.[298]

4.8.6 Analyzing "millions" of log records per second – Sentinel, Azure Data Explorer and KQL

Internally, Sentinel is built on top of Microsoft's "Azure Monitor" system and uses its "Log Analytics" functionality to store the data. **Azure Monitor's Log Analytics** system stores both data from external sources imported via Sentinel's data connectors and data created by Sentinel itself during its analysis activity.[299]

Log Analytics, in turn, is built on top of "Azure Data Explorer",[300] which can be understood as a "data lake" solution, i.e. a centralized data repository that combines and stores different kinds of data across an organization in order to make it available for analysis.[301] Microsoft describes **Azure Data Explorer** as a "big data analytics platform that makes it easy to analyze high volumes of data in near real time".[302] It is a multi-purpose system that can "ingest and analyze petabytes of data in real time".[303] It can analyze "structured, semi-structured, and unstructured data" in "different formats and structures, flowing from various pipelines and sources" and "extract key insights, spot patterns and trends".[304] Microsoft emphasizes that the system is specifically "useful for log analytics".[305] It can "analyze petabytes of log streams in real time",[306] process data on "millions of events per second" and "query petabytes of data, with results returned within milliseconds to seconds".[307] Organizations can ingest data via "queued ingestion or streaming ingestion",[308] i.e. they can either import data in regular intervals or continuously import records as they are created by a source system. Microsoft's Kusto Query Language (KQL) is part of Azure Data Explorer[309] and can be used to query any data available in Data Explorer. Microsoft systems such as Log Analytics, Sentinel and Defender for Endpoint extend the functionality provided by Azure Data Explorer and KQL.[310]

Consequently, Microsoft Sentinel can process **millions of log and activity records per second** and quickly analyze large volumes of data in real time. Almost any functionality in Sentinel, including pre-built analyses and reports, is based on KQL.[311] To search log data across long time spans, Sentinel provides "search jobs" that allow organizations

---

[296] https://learn.microsoft.com/en-us/azure/sentinel/kusto-overview [15.4.2024]
[297] https://learn.microsoft.com/en-us/azure/sentinel/hunting [15.4.2024]
[298] https://learn.microsoft.com/en-us/azure/sentinel/bring-your-own-ml [15.4.2024]
[299] https://learn.microsoft.com/en-us/azure/sentinel/kusto-overview [15.4.2024]
[300] https://learn.microsoft.com/en-us/azure/azure-monitor/logs/log-analytics-overview [15.4.2024]
[301] See e.g.: https://en.wikipedia.org/wiki/Data_lake
[302] https://learn.microsoft.com/en-us/azure/data-explorer/data-explorer-overview [15.4.2024]
[303] https://azure.microsoft.com/en-us/products/data-explorer [15.4.2024]
[304] https://learn.microsoft.com/en-us/azure/data-explorer/data-explorer-overview [15.4.2024]
[305] Ibid.
[306] https://azure.microsoft.com/en-us/products/data-explorer [15.4.2024]
[307] https://learn.microsoft.com/en-us/azure/data-explorer/data-explorer-overview [15.4.2024]
[308] https://learn.microsoft.com/en-us/azure/data-explorer/data-explorer-overview [15.4.2024]
[309] Ibid.
[310] Ibid.
[311] https://learn.microsoft.com/en-us/azure/sentinel/kusto-overview [15.4.2024]

to "find specific events in logs up to seven years ago".[312] Organizations can also restore log data that was already archived.[313] As such, Sentinel can **analyze and search up to seven years of log and activity data**.

4.9 Employee surveillance and making it transparent with the "audit log"

Purview provides **auditing functionality** that can act both as another employee surveillance tool and as a means to help organizations monitor surveillance activities or even allow employee representatives to monitor misuse. Auditing should "help organizations effectively respond to security events, forensic investigations, internal investigations, and compliance obligations", according to Microsoft. The audit log, which is activated by default, records and retains information about "thousands" of different user and admin activities performed in various Microsoft systems. It provides "visibility into the activities performed across […] Microsoft 365" for cybersecurity staff, IT admins, insider risk teams and compliance investigators.[314] Figure 22 (right) shows how the audit log can be searched for activity records according to criteria such as user/employee name, activity name and file name.

The audit log records, for example, activities in Microsoft Exchange (e.g. sending an email message, copying an email message to another folder), Teams (e.g. logging in, information about meeting participants including the time they joined and left the meeting, sending a message with an URL link), SharePoint and OneDrive (e.g. accessing, creating, modifying, moving, renaming, deleting, downloading or uploading files), Microsoft Project (e.g. creating or accessing a project, creating or accessing a task) and CoPilot (interacting with it, e.g. entering a prompt).[315]

Microsoft's security and risk monitoring systems are deeply integrated with the audit log. The Purview insider risk system requires the audit log to be enabled and utilizes it as a data source for activity monitoring.[316] Purview's communication compliance system stores alerts about suspicious communication activities in the audit log, including information about the employee who performed the activity. Sentinel and other SIEM systems can use the audit log as a data source for activity monitoring, for example, to analyze alerts from the communication compliance system.[317] Organizations can also utilize the log to provide **audit trails** about employee behavior and communication activity to compliance or regulatory auditors, both internally and externally.[318]

[312] https://learn.microsoft.com/en-us/azure/sentinel/search-jobs [15.4.2024]
[313] https://learn.microsoft.com/en-us/azure/sentinel/restore [15.4.2024]
[314] https://learn.microsoft.com/en-us/purview/audit-solutions-overview [30.4.2024]
[315] https://learn.microsoft.com/en-us/purview/audit-log-activities [30.4.2024]
[316] https://learn.microsoft.com/en-us/purview/insider-risk-management-configure [30.4.2024]
[317] https://learn.microsoft.com/en-us/purview/communication-compliance-siem [30.4.2024]
[318] https://learn.microsoft.com/en-us/purview/communication-compliance-reports-audits [30.4.2024]

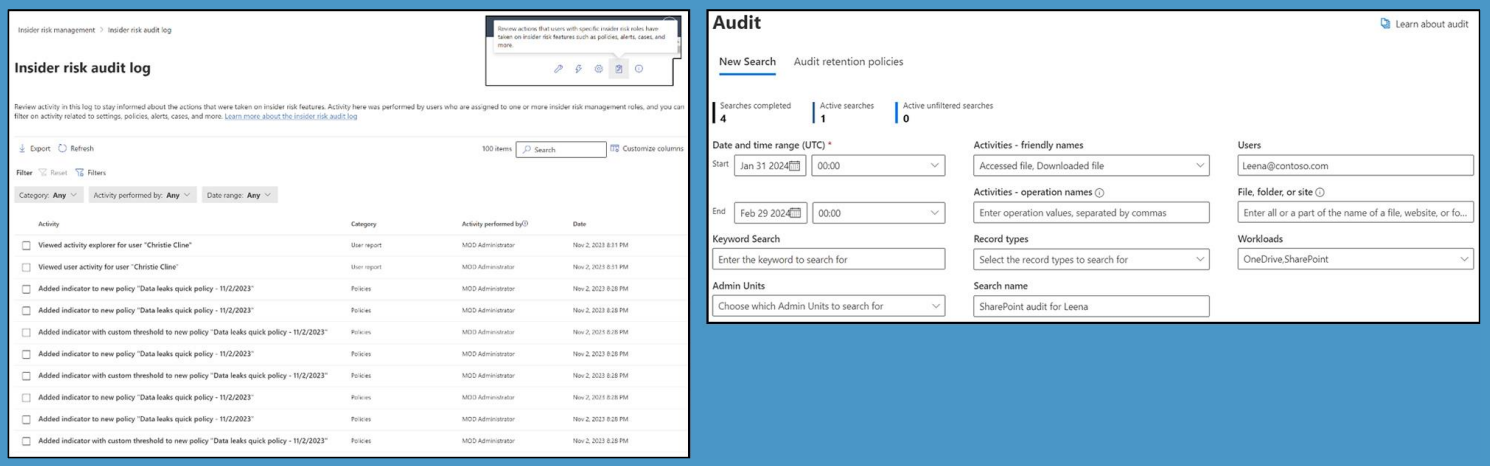

Figure 22: Insider risk audit log and searching the Microsoft 365 audit log (Microsoft Purview)[319]

By default, audit log records are retained for one year (Entra, Exchange, OneDrive, SharePoint) or 180 days (all other Microsoft systems). They can be retained for longer periods up to 10 years, which can help with "forensic or compliance investigations", "long running investigations" and responses to "regulatory, legal, and internal obligations", according to Microsoft. Audit log information can be exported via CSV and API.[320]

Monitoring employee surveillance. The audit log also records activities performed by analysts, investigators and other persons in Microsoft's security and risk profiling systems. It records, for example, activities performed in the eDiscovery system (e.g. creating and managing cases, starting searches and viewing search results, viewing and exporting documents) and in the communications compliance system (e.g. updating policies).[321] Both Microsoft Sentinel and Purview's insider risk system provide extra auditing functionality that is separate from Microsoft's "unified" audit log described above. Sentinel's audit logs record activities such as creating, updating and deleting data connectors, workbooks, watchlists and alert rules. They can, for example, provide information about the persons who performed the most Sentinel queries in the previous week.[322] The insider risk audit log records information about alerts and cases including personal data on employees, information on analysis and investigation activities and information on changes in risk policies and settings.[323] The example shown in Figure 22 (left) demonstrates how the audit log displays information about activities that involve changing risk policies and viewing activity records about suspicious employees. The insider risk audit log cannot be disabled.[324] While audit logs represent additional monitoring, they can **potentially prevent misuse**, especially when employee representatives have access to them.

[319] Figures © Microsoft, Mark Schuijt, AdminDroid. The figures serve as basis for the discussion of the corporate practices examined in this study. Sources: https://www.linkedin.com/pulse/microsoft-purview-insider-risk-management-vs-cis-bio-deep-schuijt-k7bve/, https://blog.admindroid.com/unified-audit-log-a-guide-to-track-office-365-activities/ [30.4.2024]

[320] https://learn.microsoft.com/en-us/purview/audit-solutions-overview [30.4.2024]

[321] https://learn.microsoft.com/en-us/purview/audit-log-activities, https://learn.microsoft.com/en-us/purview/ediscovery-search-for-activities-in-the-audit-log [30.4.2024]

[322] https://learn.microsoft.com/en-us/azure/sentinel/audit-sentinel-data [30.4.2024]

[323] https://learn.microsoft.com/en-us/purview/insider-risk-management-audit-log [30.4.2024]

[324] Ibid.

4.10 Employee privacy, data protection and other safeguards?

As this case study shows, Microsoft's cybersecurity and risk profiling systems potentially process extensive personal data about employee behavior and communication. In many cases, this involves singling out suspicious employees, ranking them by risk level and investigating their past activity at the individual level. While employers can customize data sources and processing in many ways, Microsoft provides a range of intrusive risk policies out of the box,[325] incentivizes organizations to implement intrusive surveillance via its "compliance center" software[326] and sometimes recommends that organizations implement the more intrusive options.[327]

Pseudonymization in Microsoft Purview. The insider risk and communication compliance systems offer to replace employee names, email addresses and other personal data displayed in the user interface with pseudonyms such as "AnonyIS8-978". While several example screens shown in sections 4.1 to 4.8 display employee names alongside risk scores, rankings and detailed behavioral records, pseudonymization is now turned on by default for user roles such as "insider risk management analysts", "insider risk management investigators" and "communication compliance analysts", according to Microsoft.[328] Pseudonymization certainly offers some protection from arbitrary spying on employees by unauthorized staff. Nevertheless, employee names can still be accessed by those who have the permission to access them. Employers can decide to turn pseudonymization on or off for certain user roles.[329] Figure 11 (bottom left) shows how pseudonymization can be turned on or off in the user interface.

Pseudonymization in Microsoft Sentinel. Sentinel also often displays employee names, email addresses or other personal data in the user interface, as shown in section 4.8. As of April 2024, the Sentinel online documentation does not address data protection.[330] Apparently, the system does not provide functionality to replace personal data with pseudonyms in the user interface.[331] Both activity log files ingested into Sentinel and entity information analyzed by the system will often contain device IDs, user IDs or usernames referring to employees,[332] and as such, personal data on employees. The Advanced Security Information Model (ASIM), which describes a standard for activity log records promoted by Microsoft, mentions Windows user IDs, Linux user IDs, Entra user IDs, Microsoft 365 user IDs, Amazon AWS user IDs and Salesforce user IDs.[333] Entity usernames analyzed by Sentinel can be "anonymized", which likely refers to pseudonymization rather than anonymization. They are, however, not "anonymized" by default, according to the documentation.[334] On the contrary, Microsoft prominently suggests that "data about […] user accounts, including the user identification and privileges, are crucial for the analysts in the process of an investigation". Consequently, organizations that use Sentinel's UEBA functionality can, for example, "correlate security events with the IdentityInfo table", which "synchronizes" with account and identity information stored

[325] E.g. utilizing data on "offensive language", "risky browser usage" or "disgruntled" employees to assess "insider risk"; see section 4.1
[326] See section 4.1.1
[327] E.g. Microsoft recommends to monitor "all" employees in an organization at least for "harassment or discrimination detection", see section 4.2
[328] https://learn.microsoft.com/en-us/purview/insider-risk-solution-privacy [7.3.2024]
[329] Ibid.
[330] The author's Google searches for the terms "data protection", "personal data", "gdpr" and "dsgvo" did not show any meaningful results. Google search phrases used: site:learn.microsoft.com/en-us/azure/sentinel "data protection"; site:learn.microsoft.com/en-us/azure/sentinel "personal data"; site:learn.microsoft.com/en-us/azure/sentinel "gdpr"; site:learn.microsoft.com/de-de/azure/sentinel "dsgvo" [25.4.2024]
[331] Google search phrases used: site:learn.microsoft.com/en-us/azure/sentinel "pseudonymization"; site:learn.microsoft.com/en-us/azure/sentinel "pseudonyms"; site:learn.microsoft.com/en-us/azure/sentinel "pseudonymize", site:learn.microsoft.com/en-us/azure/sentinel "pseudonymized"
[332] https://learn.microsoft.com/en-us/azure/sentinel/normalization-about-schemas#the-user-entity, https://learn.microsoft.com/en-us/azure/sentinel/normalization-common-fields, https://learn.microsoft.com/en-us/azure/azure-monitor/logs/personal-data-mgmt, https://learn.microsoft.com/en-us/azure/sentinel/entities-reference [25.4.2024]
[333] https://learn.microsoft.com/en-us/azure/sentinel/normalization-about-schemas#the-user-entity [25.4.2024]
[334] Ibid.

in Microsoft Entra, and thus adds employee information to the analysis of logs that do not directly contain this information in the first place.[335]

The online documentation for Microsoft's "Defender for Cloud Apps" system,[336] which monitors activity in cloud-based applications and can share data with Purview and Sentinel,[337] provides more solid information about pseudonymization in log data. The system can replace usernames with "encrypted" usernames, i.e. with pseudonyms.[338] However, "admins" can always "resolve the real username" for a "specific security investigation" such as "a security breach or suspicious user activity". If "an admin has a reason to suspect a specific user, they can also look up the encrypted username of a known username, and then start investigating using the encrypted username".[339] As such, this kind of pseudonymization only offers very limited protection for employees. The "Defender for Cloud Apps" documentation specifically mentions that organizations can limit monitoring to certain groups rather than using it "for all users". Microsoft explains that "compliance regulations" may require organizations to "not monitor users from certain countries/regions". For example, they could "only monitor US-based employees" and "avoid showing any activities" for employees "based in Germany".[340]

Roles and permissions. Microsoft's security and risk profiling systems use a wide range of "permissions" and "roles" that determine who can access or modify what kind of data. Sentinel provides a number of pre-configured roles. A "Sentinel Reader", for example, can view incidents, workbooks and other Sentinel data. A "Sentinel Responder" can additionally manage and modify incident data. A "Sentinel Contributor" can install, create and modify components such as workbooks and analytics rules. Microsoft recommends assigning "security analysts" the "Sentinel Responder" role and "security engineers" the "Sentinel Contributor" role. Higher-level roles such as the "Azure Reader" or "Log Analytics Reader" also grant access to some Sentinel information. Installing additional data sources may require the "Global Administrator" or "Security Administrator" roles. Organizations can create custom roles or customize the permissions linked to a role.[341]

Microsoft Purview also offers a range of pre-configured roles and "role groups", the latter of which refers to sets of roles that can be assigned to a person. The communication compliance system provides separate role groups that allow persons to view information ("Communication Compliance Viewers"), investigate policy matches and view message metadata ("Communication Compliance Analysts") or additionally view message content ("Communication Compliance Investigators"). "Communication Compliance Administrators" can create and modify policies and settings. Purview's insider risk system provides similar role groups that allow persons to access alerts and cases ("Insider Risk Management Analysts"), additionally access the "content explorer" for all cases ("Insider Risk Management Investigators") or create and modify policies and settings ("Insider Risk Management Admins").[342] Purview's eDiscovery and data loss prevention systems also provide pre-configured roles that allow persons to put mailboxes and other resources "on hold", search for information, review it, access data loss prevention alerts, investigate activities and contents or configure policies and settings.[343]

---

[335] https://learn.microsoft.com/en-us/azure/sentinel/investigate-with-ueba [25.4.2024]
[336] https://learn.microsoft.com/en-us/defender-cloud-apps/editions-cloud-app-security-o365 [25.4.2024]
[337] https://learn.microsoft.com/en-us/defender-cloud-apps/cas-compliance-trust [25.4.2024]
[338] Ibid.
[339] https://learn.microsoft.com/en-us/defender-cloud-apps/cloud-discovery-anonymizer [25.4.2024]
[340] https://learn.microsoft.com/en-us/defender-cloud-apps/scoped-deployment [25.4.2024]
[341] https://learn.microsoft.com/en-us/azure/sentinel/roles [25.4.2024]
[342] https://learn.microsoft.com/en-us/defender-office-365/scc-permissions [25.4.2024]
[343] Ibid.

Persons who are assigned to the higher-level role group "Insider Risk Management" can access and modify almost everything in the system. Microsoft explains that organizations can "use this role group to manage insider risk management for your organization in a single group", which is "the easiest way to quickly get started with insider risk management and is a good fit for organizations that don't need separate permissions defined for separate groups of users". By suggesting to add "all user accounts for designated administrators, analysts, and investigators" to a single group, Microsoft undermines its own security and privacy model. Other higher-level role groups allow persons to access and modify almost everything in the communication compliance, eDiscovery or data loss prevention system. Yet other role groups provide access to functionality across different Purview systems (e.g. "Organization Management") or assign roles to persons ("Purview Administrator"). The "Global Reader" role group provides read-only access to all information in Purview, including reports, alerts and settings.[344]

Exporting personal data. Both Sentinel and Purview allow information, including personal data on employees, to be exported to external files or systems via CSV or API access. This includes log data and alerts from Sentinel,[345] alerts from Purview's insider risk system,[346] alerts and message contents from Purview's communication compliance system[347] and eDiscovery contents.[348] Obviously, Microsoft's role and permission protections no longer apply as soon as data is exported to files or third-party systems.

Audit logs. As detailed in section 4.9, Microsoft's audit logs can act both as additional employee surveillance tools and as potential safeguards that prevent misuse of the company's powerful security and risk profiling systems. Audit logs can record information about how security analysts, compliance investigators and other persons make use of Purview and Sentinel, from accessing information about employees who were assessed as suspicious to changes in settings and risk policies. It is, however, questionable whether employees can trust an organization to audit itself. To ensure robust protection from an employee perspective, auditing must be performed either by departments or external parties who are independent from the organization's management or by employee representatives.

Microsoft's data protection claims. Microsoft generally emphasizes that it "practices privacy by design and privacy by default in its engineering and business functions" and that it "performs comprehensive privacy reviews on data processing operations that have the potential to cause impacts to the rights and freedoms of data subjects". It claims to "review the design and implementation of services to ensure that personal data is processed in a respectful manner that accords with international law, user expectations, and our express commitments". Microsoft also states that organizations that use the company's systems "are required to prepare a Data Protection Impact Assessment (DPIA) for processing operations" that are "likely to result in a high risk to the rights and freedoms of natural persons" under the GDPR. The company also claims, however, that there is "nothing inherent in Microsoft products and services that need the creation of a DPIA. Rather, it depends on the details of [the organization's] Microsoft configuration".[349] For its cloud-based Azure environment, Microsoft states that "Azure services are not designed to perform processing on which decisions are based that produce legal or similarly significant effects on individuals" and they are also "not designed to process special categories of personal data on a large scale". However, "because

[344] Ibid.

[345] https://techcommunity.microsoft.com/t5/microsoft-sentinel-blog/export-historical-log-data-from-microsoft-sentinel/ba-p/3413418, https://learn.microsoft.com/en-us/azure/sentinel/kusto-overview, https://learn.microsoft.com/en-us/azure/data-explorer/kusto/management/data-export/, https://techcommunity.microsoft.com/t5/microsoft-sentinel-blog/sending-enriched-microsoft-sentinel-alerts-to-3rd-party-siem-and/ba-p/1456976 [30.4.2024]

[346] https://learn.microsoft.com/en-us/purview/insider-risk-management-settings-dlp-sync [30.4.2024]

[347] https://learn.microsoft.com/en-us/purview/communication-compliance-siem, https://learn.microsoft.com/en-us/purview/communication-compliance-reports-audits, https://learn.microsoft.com/en-us/purview/communication-compliance-investigate-remediate [30.4.2024]

[348] https://learn.microsoft.com/en-us/purview/ediscovery [30.4.2014]

[349] https://learn.microsoft.com/en-us/compliance/regulatory/gdpr [25.4.2024]

Azure is a highly customizable service, a data controller could potentially configure Azure services to be used for such processing. Controllers should make this determination based on their usage of Azure".[350] In light of the powerful surveillance capabilities provided by Purview and Sentinel, several of these claims are questionable.

5. Systems from other vendors – IBM and Teramind

As outlined in section 2.1, several other vendors offer software similar to the systems provided by Forcepoint and Microsoft. This section briefly examines software provided by IBM and Teramind.

Security and risk profiling systems provided by large enterprise software vendors such as Microsoft and IBM provide intrusive surveillance capabilities that can be used for a many different purposes. While they promote them as solutions to very different problems, the purposes generally remain within the boundaries of cybersecurity, risk management and compliance. Smaller software vendors such as Teramind provide systems that go beyond that.

IBM's cybersecurity system QRadar offers SIEM, UEBA and insider threat detection functionality.[351] It can monitor "millions" of activity records from event logs, employee devices and network data "in near real time" across 450 data sources and 370 applications and analyze them "against historical data to uncover known and unknown threats".[352] The UEBA module, which IBM refers to as "user behavior analytics" (UBA), analyzes "behavioral patterns" and builds "risk profiles" about users and employees in order to understand "normal behavior" and detect "anomalous behavior". It promises to detect different kinds of suspicious activities ranging from "identity theft, hacking, phishing or malware" to "insider threats",[353] data theft and non-compliance with "enterprise, industry, and regulatory" policies.[354] Similar to the systems provided by Forcepoint and Microsoft, it calculates "risk scores" for suspicious users and employees. Behavioral profiling can be based on rules or AI-based models and can include activities such as visiting "harmful or compromised" websites.[355] A promotional video provided by IBM shows how QRadar displays detailed information on a named employee who was assessed as a risk because of activities such as visits to "gambling" and "lifestyle" websites, among others.[356] In addition to data from applications and "endpoint" devices, it can analyze user behavior based on monitoring network traffic (e.g. via "proxies, firewalls, IPS, and VPNs").[357] For example, it can detect "sensitive content" or "negative sentiment" in communication activity, including in email content.[358] The system can build "unified user identities" based on data on users and employees from different sources in order to "combine risk and traffic data across different usernames".[359] Furthermore, QRadar provides functionality for searching large amounts of activity log data[360] and conducting "forensic" investigations for purposes such as "network security, insider analysis, fraud and abuse, and evidence-gathering".[361]

---

[350] https://learn.microsoft.com/en-us/compliance/regulatory/gdpr-dpia-azure [25.4.2024]
[351] https://www.ibm.com/qradar, https://www.ibm.com/products/qradar-siem/user-behavior-analytics [5.5.2024]
[352] https://www.ibm.com/products/qradar-siem/advanced-threat-detection [5.5.2024]
[353] https://www.ibm.com/products/qradar-siem/user-behavior-analytics [5.5.2024]
[354] https://www.ibm.com/docs/en/qsip/7.5?topic=insights-qradar-network-use-cases [5.5.2024]
[355] https://www.ibm.com/products/qradar-siem/user-behavior-analytics [5.5.2024]
[356] Video „IBM Security QRadar SIEM – User Behavior Analytics", min 1:25: https://www.ibm.com/products/qradar-siem/user-behavior-analytics [5.5.2024]
[357] https://www.ibm.com/docs/en/qradar-common?topic=app-qradar-user-behavior-analytics [5.5.2024]
[358] https://www.ibm.com/docs/en/qsip/7.5?topic=insights-qradar-network-use-cases [5.5.2024]
[359] https://www.ibm.com/products/qradar-siem/user-behavior-analytics [5.5.2024]
[360] https://www.ibm.com/products/qradar-log-insights [5.5.2024]
[361] https://www.ibm.com/docs/en/qsip/7.5?topic=forensics-security-investigations [5.5.2024]

IBM also offers systems for **communication monitoring** in the financial sector. Its "Surveillance Insight for Financial Services" product uses profiling, "behavior analysis" and "anomaly detection" to detect "non-compliant behavior" and "misconduct"[362], such as insider trading and "market abuse activity", in financial organizations.[363] The system's "Electronic Communication Surveillance Analytics" module monitors and profiles the contents of employee communication via email, chat and voice calls.[364] The analysis of voice calls is based on automated transcripts.[365] The system promises to detect behavioral "anomalies" such as "unusual communication timings", the "intent to use insider information", efforts to "recruit co-conspirators", "confidential anomaly", "negative sentiment anomaly", "anger anomaly" and "sad anomaly".[366] IBM claims to assess "emotions and sentiment" in conversations based on its "emotion detection library", which can detect "anger", "disgust", "joy", "sadness" and "fear", according to the documentation.[367] IBM's "Financial Crimes Insight for Conduct Surveillance" product provides similar functionality[368] and also includes emotion detection.[369] It "monitors employee activity by ingesting and analyzing the immense volume of data derived from [a] diverse set of new channels and source[s], including emails, social media, chat transcripts, voice transcripts and customer complaints".[370]

Teramind – combining risk and productivity monitoring. The US-based software vendor Teramind offers an employee surveillance system that provides "insider threat prevention", "insider fraud detection", DLP, UEBA and "compliance management", according to the company.[371] The market research firm Gartner lists Teramind in its market guide for insider risk management solutions (Gartner 2020). However, the same system can also be used for "employee productivity monitoring", for example, to identify "high- and low-performing employees" and "unproductive employee activity".[372] Employers can use it to "capture, analyze and control user desktop activity for any use case", according to Teramind.[373] Based on software installed on the computers and devices of employees, the system monitors application use, website visits, email and message communication, online meetings and file activity. It classifies activities such as application uses and website visits into categories like "productive" or "unproductive", calculates risk scores for employees and provides a wide range of employee rankings and details about their activities. The system can even record keyboard activity and allows organizations to search for employees who entered certain keystrokes. It can record all screen activity on employees' computers and provides access to "video-quality session recordings" both in real-time and for certain points in time in the past. Not least, the system provides full remote access to employee computers.[374]

Consequently, the system can be considered fully-fledged "spyware" or even a "rootkit"[375] with functionality that is typically included in computer viruses and other malicious software. According to a Teramind representative indirectly quoted by the Washington Post, Teramind "cannot implement safeguards without significantly hindering

[362] https://www.ibm.com/docs/en/siffs/2.0.3?topic=using-surveillance-insight-financial-services [14.5.2024]
[363] https://www.ibm.com/docs/en/siffs/2.0.3?topic=services-trade-surveillance-analytics [14.5.2024]
[364] https://www.ibm.com/docs/en/siffs/2.0.3?topic=services-electronic-communication-surveillance-analytics [14.5.2024]
[365] https://www.ibm.com/docs/en/siffs/2.0.3?topic=services-voice-surveillance-analytics [14.5.2024]
[366] https://www.ibm.com/docs/en/siffs/2.0.3?topic=services-electronic-communication-surveillance-analytics [14.5.2024]
[367] https://www.ibm.com/docs/en/siffs/2.0.3?topic=learning-emotion-detection-library [14.5.2024]
[368] https://www.ibm.com/docs/en/fci/6.5.1?topic=surveillance-introducing-financial-crimes-insight-conduct [14.5.2024]
[369] https://www.ibm.com/docs/en/fci/6.5.1?topic=communication-discovery [14.5.2024]
[370] https://www.ibm.com/docs/en/fci/6.5.1?topic=insight-what-is-financial-crimes [14.5.2024]
[371] https://www.teramind.co/ [5.5.2024]
[372] https://www.teramind.co/solutions/workforce-productivity-optimization, https://www.teramind.co/features/unproductive-work-time-analysis [5.5.2024]
[373] https://www.teramind.co/ [5.5.2024]
[374] Ibid.
[375] See e.g. https://en.wikipedia.org/wiki/Rootkit

the software's capability".[376] Because of its intrusive multi-purpose functionality and the lack of adequate safeguards, it is unlikely that European employers can use the system lawfully under the GDPR. It is however concerning that Teramind's software can easily be bought by the smallest businesses starting from €10.50 per employee. The fact that prices are indicated in Euros suggests that Teramind is targeting European markets.[377]

[376] https://www.washingtonpost.com/technology/2021/08/20/work-from-home-computer-monitoring/
[377] https://www.teramind.co/product/price [5.5.2024]

6. Summary of data practices that affect employees

Any organization will take measures to protect itself from cyberattacks, data loss and other information security threats. This is not optional, and it is in several ways mandated by law — in Europe, for example, by the GDPR and the new NIS2 directive.[378] Furthermore, any organization will aim to protect itself from intentional data leaks, fraud and unlawful employee behavior and enforce compliance with regulatory obligations and organizational policies.

This case study examines cybersecurity and risk profiling software that promises to help organizations meet these challenges. It investigates systems for "security information and event management" (SIEM), "user and entity behavior analytics" (UEBA), insider risk management, communication monitoring and eDiscovery provided by Forcepoint, Microsoft and other vendors. These systems typically process extensive personal data on employees, allow organizations to monitor their employees' every step and single them out for further investigation. Depending on how these systems are implemented, they can enable far-reaching employee surveillance, both across entire organizations and targeted at individuals. This case study focuses on the potential ramifications for employees.

The following sections present an overview of data practices identified and documented in this investigation.

6.1 Forcepoint

The insider risk and UEBA systems provided by **Forcepoint**,[379] a US cybersecurity vendor closely associated with the defense and intelligence sector, monitor a wide range of employee behaviors and communication activities based on data collected from different sources including employee devices, corporate networks and enterprise software systems from vendors such as Microsoft, Salesforce, Cisco and Workday. As detailed in section 3 and summarized in Table 5, Forcepoint's security and risk profiling systems can monitor how employees use files and applications, the websites they visit and their web searches, email and chat conversations, voice call transcripts, their screen, keyboard and clipboard activity, logins, physical access to buildings and rooms, geolocation data, performance reviews from HR systems and external data (e.g. on "financial distress"). Forcepoint aims to help organizations detect "anomalous" or otherwise suspicious behavior based on profiling in order to identify cyberattacks and employees who are a considered a risk, whether by carelessness, negligence or intention. Based on digital profiling, it continuously calculates risk scores for employees, singles out those who are assessed as suspicious, ranks them by risk and raises alerts. To identify "anomalous" behaviors, it analyzes behavioral data about many employees.

Forcepoint promises to assess whether employees are in financial distress, show "decreased productivity" or intend to leave the job, how they communicate with colleagues and whether they access "obscene" content or exhibit "negative sentiment" in their communications. Its monitoring technology can also detect custom keywords in communication activity and visits to websites that are categorized as "legal liability" or "productivity loss". Furthermore, it promises to detect "illicit workplace behavior" such as "corporate espionage" and "whistle-blowing" by employees who are "in contact with media domains" and show "signs of willingness to leak to leak company information". Companies in the financial sector can detect "insider trading" and "market manipulation". Organizations can then further investigate suspicious employees and their past behavior, including their website visits, application and file usage, email and chat contents, keyboard activity (e.g. typing, copying to the clipboard), screen recording and even

[378] See e.g. https://www.enisa.europa.eu/topics/cybersecurity-policy

[379] Note: The research in this case study refers to products offered by Forcepoint up until late 2023 (see section 3)

the time they spend on different activities. Forcepoint enables "an 'over-the-shoulder' view of the end-user's workstation", according to the company. It provides "forensic evidence" for "investigations, prosecution, and compliance" and "unparalleled visibility into suspicious behaviors", which is "admissible in a court of law". Table 5 summarizes how Forcepoint's software processes personal data on employees, according to the findings in section 3.

System, module	Analyzed data categories, data sources	Personal data processing on employees	Examples of purposes and detected behaviors
Forcepoint UEBA and insider risk systems Monitors employee behavior and communication content to detect cyberattacks, "insider threats" and undesirable behavior (section 3)	• Data on login activity, file activity (e.g. view, create, modify, delete, move, download, share, copy to clipboard), application activity, web browser activity (e.g. websites visited, Google searches, webmail usage), network activity (proxy, VPN), printing activity • Extended device monitoring data (screen recording, keyboard activity, clipboard usage) • Communication contents and metadata (email, chat, SMS, voice call transcripts) • Badging data on building and room access • "Geolocation/GPS" data • HR information, e.g. performance reviews, promotions, compensations • External data, e.g. criminal records, financial distress data • Data from other Forcepoint systems, e.g. endpoint device monitoring, DLP • Data from third-party cybersecurity (e.g. SIEM, DLP), networking (e.g. proxy) and endpoint device monitoring systems (e.g. Veriato, Digital Guardian) • Data from third-party enterprise software systems, e.g. Microsoft (Windows, Active Directory, Exchange, Office 365, Skype), Salesforce (CRM, Slack), SAP (e.g. Concur), Cisco (network infrastructure), Workday (HR)	• Process and analyze extensive personal data on employee behavior and communication activity • Monitor and profile employee behavior and employee communication for different purposes over time • Detect "anomalous" or otherwise suspicious behavior based on risk "scenarios", "models" and behavioral profiles on employees and their peers • Detect keywords, phrases and "sentiment" in communication activity • Continuously calculate risk scores for activities and employees • Single out employees whose behavior was detected as suspicious, rank them by risk and raise alerts • Share alerts with other systems • Investigate past employee activity in detail, e.g. website visits, application and file usage, keyboard activity (e.g. "typing", "copying"), screen recordings, email and chat contents • Investigate the time employees spent on different activities • Monitor and profile communication with persons outside the organization	• Detect cyberattacks, compromised user accounts and employees that are considered a threat to an organization (e.g. "careless" and "negligent" employees, "data thieves", "disgruntled employees" who had a "huge fight with the boss", "saboteurs" or "media leakers") • Detect employees who "are behaving different than usual" and thus "may need to be investigated" • Detect "off hour" and "weekend" activity • Detect "data exfiltration", "data loss", "anomalous movement of data", "anomalous interactions with files", "anomalous number of distinct file shares accessed", "activity conducive with searching for data", "employees who are leaking sensitive information" • Detect "suspicious users", "abnormal authentication activity", "interactions with core system files", "people researching ways to commit suspicious actions", "abnormal physical access" • Detect "compromised user accounts" and employees who fell victim to phishing or malware installation • Detect "negative workplace behavior" such as "obscene" web searches or website visits or "negative sentiment" and "improper discussions" within communications • Detect employees who show "decreased productivity", who are "not interacting with core company assets" and show "corporate disengagement", who are communicating with fewer colleagues than usual, who are in "financial distress" or who plan to leave the job • Detect "workplace violence" and "sexual harassment" • Detect "illicit workplace behavior" such as "corporate espionage", "whistleblowing" (persons who are "in contact with media domains" showing "signs of willingness to leak company information"), "clearance evasion" ("people researching ways to omit security clearance information or ways to deceive a polygraph") • Detect "insider trading" and "market manipulation" • Detect visits to websites that are classified as malware, "legal liability" (e.g. adult, gambling, file sharing, weapons) or "productivity loss" (e.g. social media, drugs, health, abortion, religion, worker organizations) • Provide "forensic evidence" for "investigations, prosecution, and compliance" and "unparalleled visibility into suspicious behaviors", enable "an 'over-the-shoulder' view of the end-user's workstation", allowing for "attribution as well as showing employee intent", which is "admissible in a court of law"

Table 5: Data practices identified in this investigation: Forcepoint's UEBA and insider threat systems

6.2 Microsoft

The enterprise software giant **Microsoft** offers cybersecurity and risk profiling software that is deeply integrated with its cloud-based Microsoft 365 system and can be integrated with other enterprise software. Microsoft's SIEM system "Sentinel" collects and analyzes data from many sources for cybersecurity purposes and includes UEBA functionality. Microsoft's "Purview" system addresses data security, risk management and compliance. **Microsoft Purview** provides software systems for insider risk management, communication monitoring ("Communication Compliance") and for searching employee information ("eDiscovery"), among others. Table 6 summarizes how these systems process personal data on employees, according to the findings in section 4.

System, module	Analyzed data categories, data sources	Personal data processing on employees	Examples of purposes and detected behaviors
Microsoft Purview "Communication Compliance" Monitors and scans communication and file contents for different purposes (sections 4.2, 4.7) *Can share alerts about employees with "Insider Risk Management", Microsoft Sentinel and other systems*	• Communication and document contents (e.g. emails, chat messages, documents and files, text in images, recorded video meetings and transcripts, SMS, voice calls and transcripts, Copilot prompts and responses) • Data from Microsoft systems, e.g. Microsoft 365, Exchange, Teams, Viva Engage and Microsoft Copilot • Data from third-party software systems for communication (e.g. Zoom, Slack, Cisco Webex, RingCentral) and financial messaging (e.g. Instant Bloomberg, ICE Chat) • Data from messengers via on-device access, e.g. WhatsApp/Signal via TeleMessage • Data on voice calls and SMS from standard mobile phones, e.g. via TeleMessage on-device access or mobile carrier partnerships • Data from other systems via custom integrations, e.g. SQL databases, CSV, XML, EML	• Process and analyze extensive personal data on communication and file content • Monitor, profile and scan employee communication for different purposes • Detect suspicious employee behavior according to customizable keywords and AI-based content classifiers • Single out employees whose communication was detected as suspicious • Block communication, notify employees, notify managers about alerts, create tasks, trigger automated workflows • Share alerts with other systems; export alerts and suspicious content • Investigate past communication contents and metadata • Monitor and profile communication with persons outside the organization • Create aggregate reports on suspicious behavior in the organization	• Detect "inappropriate messages" to "minimize communication risks"; identify "potential legal exposure and risk" and non-compliance with regulatory standards and corporate policies (e.g. "acceptable use, ethical standards") • Detect "inappropriate" text, "profanity", "expressions that embarrass most people", "offensive language", hate, violence, self-harm, "sexual" content, "inappropriate" images, "adult, racy and gory" images • Detect "threatening language", "threats to commit violence or do physical harm or damage to a person or property", "targeted harassment", "discrimination" • Detect bribery, gift exchanges, conflicts of interest, money laundering, insider trading, stock manipulation • Detect "workplace collusion", "secretive actions such as concealing information or covering instances of a private conversation, interaction, or information" • Detect corporate sabotage, unauthorized disclosure of confidential information, insider threats • Detect any pattern based on custom classifiers • Classify sentiment of suspicious contents
Microsoft Purview "Insider Risk Management" Monitors employee behavior to detect compromised accounts and undesirable employee behavior (sections 4.1, 4.3-4.7) *Includes "Communication Compliance" functionality; can receive data from and share alerts with Microsoft Sentinel and other systems; can share user risk scores with Purview DLP; can trigger eDiscovery investigations*	• Log data about logins, file activity (e.g. view, create, modify, delete, move, download, share, copy to clipboard), application activity, web browsing activity, printing activity, USB activity, communication activity (e.g. email, chat), meetings • Communication contents and metadata • "Forensic" device monitoring data (user interactions and screen activity) • User account data and "employee profile data" (e.g. name, job title, department, roles) • Employee performance data from the HR system (e.g. performance ratings, performance improvement plans, terminations) • Badging data on building and room access • Any other behavioral data via custom integrations (e.g. network and firewall activity, Dropbox and Salesforce activity) • Data from Microsoft enterprise systems (Microsoft 365, Exchange, SharePoint, Teams, OneDrive, Entra, Active Directory, Defender for Endpoint, Defender for Cloud Apps, Sentinel, InTune) • Data from devices and browsers monitored via Microsoft technology (e.g. Windows, macOS, Android, Edge, Chrome) • Data from Purview "Communication Compliance" and Microsoft Sentinel • Data from third-party systems (e.g. Dropbox, Salesforce, SIEM systems like Splunk)	• Process and analyze extensive personal data on employee behavior • Monitor and profile employee behavior across the organization over time • Detect suspicious employee behavior according to customizable risk policies • Detect suspicious communication content via "Communication Compliance" • Detect behavior that is "unusual" based on profiling behavior across peers and the entire organization • Detect suspicious sequences of seemingly benign activities over time • Continuously calculate risk scores for activities and employees • Single out employees whose behavior was detected as suspicious, rank them by risk and raise alerts • Notify managers about alerts, create tasks, trigger automated workflows • Share alerts and employee risk scores with other systems • Investigate suspicious employee activity in detail (e.g. files accessed, websites visited, email/message contents) • Initiate eDiscovery investigations • Process personal data on persons outside the organization • Create aggregate reports on suspicious behavior in the organization	• Detect "illegal, inappropriate, unauthorized, or unethical behavior and actions" carried out by users and employees • Detect security policy violations, data misuse, data theft, IP theft, data leaks, violations of "confidentiality obligation[s] during departure", "data access during remote work", attempts to "bypass security controls" • Detect the "misuse of patient data" by employees because of a "lack of awareness, negligence, or fraud" • Detect activities like "reading sensitive info", "copying sensitive […] content to the clipboard", physical access to "sensitive assets", "sending sensitive info in a Teams message", the use of "unwanted software", "unacceptable web usage", the use of "offensive language" in emails • Detect "risky browser usage", e.g. visits to "inappropriate or unacceptable" websites, which may "elevate network security risks", "violate regulatory requirements", expose the employer to "legal actions" or "jeopardize current and future business operations and opportunities" • Detect "inappropriate language" indicating "workplace stress", which may lead to malicious employee behavior • Detect malicious activities by employees with a "predisposition" or "tendency" to "violate company policies", because they conducted a "policy violation" in the past • Detect malicious activities by "disgruntled" employees who "experience employment stressors" such as "performance improvement notifications, poor performance reviews, changes to job level status", being "demoted or placed on a performance improvement plan" • Identify "disgruntled" employees based on "disgruntlement indicators" or a "dedicated disgruntlement policy" • Detect malicious or inadvertent activities by "risky" users
Microsoft Purview "eDiscovery" Retrieve, archive and search employee communication and document content and compile dossiers for investigations (sections 4.4, 4.7)	• Same data categories and sources as for "Communication Compliance"	• Archive and analyze personal data on employee communication and files • Identify and retrieve content associated with certain employees or topics • Process personal data on persons outside the organization • Preserve mailboxes and other contents by putting them on "legal hold" • Export the retrieved information	• Discover "electronic information" for "internal and external investigations" and as "evidence in legal cases" involving "certain executives or other employees" • Identify "persons of interest", discover data "where it lives", find emails, documents and other items "used by people in their day-to-day work tasks", retrieve data that is "most likely relevant to the case" • Retrieve data on investigation targets and others, perform targeted or mass searches across many employees

Table 6: Data practices identified in this investigation: Microsoft Purview's risk profiling systems

As detailed in section 4, Microsoft Purview monitors a wide range of employee behaviors and communication activities based on data collected from different sources including employee devices and enterprise software systems provided by Microsoft (e.g. Microsoft 365, Exchange, Teams, OneDrive) and other vendors (e.g. Zoom, Slack, Webex, Dropbox, Salesforce, Splunk). Via custom integrations, it can analyze data from almost any source.

Purview's **communication monitoring system** scans employee conversations and document content, including emails, chat messages, files, images and automated transcripts of video meetings or voice calls, for very different purposes. It detects "inappropriate messages" to "minimize communication risks" and to identify "potential legal exposure", according to Microsoft. It also helps to ensure that employees comply with "regulatory compliance standards" and with "acceptable use, ethical standards, and other corporate policies". Based on keywords and AI-based classifiers, the system promises to detect "profanity", "offensive language", "inappropriate text", "adult, racy and gory" images, threats, harassment and discrimination, but also content that indicates corporate sabotage, money laundering, bribery, conflicts of interest, insider trading, disclosure of confidential information and "workplace collusion", i.e. "secretive actions such as concealing information or covering instances of a private conversation, interaction, or information". Each "policy" that detects a certain type of undesirable communication activity can store and monitor up to one million messages. Communication and file content that was assessed as suspicious raises an alert and can be further investigated. Optionally, inappropriate content can be removed automatically. The system can also scan text entered into Microsoft's AI chatbot CoPilot, and, via third-party integrations, it can access data on calls and SMS from mobile devices and even on encrypted messages in WhatsApp or Signal.

All functionality provided by the communication monitoring system can also be utilized in Microsoft Purview's **insider risk management system**, which analyzes additional activity log data on employee behavior from devices, corporate networks, badging systems, HR systems and other enterprise software provided by Microsoft and other vendors. In addition to employee communications via email, chat and other channels, Microsoft's insider risk system can monitor how employees use files and applications, how they print files or copy them to the clipboard, the meetings they participate in, the websites they visit, their logins and physical access to buildings, offices or conference rooms. It can access data from devices, for example, via Microsoft's anti-virus system Defender or via the company's mobile device management (MDM) software InTune. To provide "forensic evidence", the system can also record screen activity and fine-grained data on every user interaction performed on employee computers. Furthermore, it can access identity data (e.g. employee name, position, job title) and track HR information (e.g. performance reviews, demotions). Via access to other cybersecurity software or custom integrations, it can monitor almost any type of activity recorded in corporate networks or connected enterprise systems (e.g. Salesforce, Dropbox).

Based on behavioral profiling, Microsoft's insider risk system aims to help organizations detect "unusual" or otherwise suspicious activities in order to identify cyberattacks, compromised user accounts and employees who are a considered a risk, whether by carelessness, negligence or intention. More broadly, it aims to detect "illegal, inappropriate, unauthorized, or unethical behavior and actions" carried out by employees, according to Microsoft. It continuously calculates risk scores for employees, singles out those who are assessed as suspicious, ranks them by risk and raises alerts. To identify sophisticated threats, it detects not only single suspicious activities but also suspicious sequences of seemingly benign activities over time. To identify "unusual" behaviors, it analyzes and profiles behaviors across the entire organization. The system promises to detect, for example, security policy violations, attempts to bypass security controls, violations of confidentiality obligations, data misuse, data and IP theft and data leaks, for example, by monitoring employees who "read" sensitive information, copy it to the clipboard, send it via email or Teams, or physically access "sensitive assets". Indicators that contribute to the assessment of employees

as potential "insider threats" can include the use of "unwanted software", the use of "offensive language" in emails, "unacceptable web usage" or "risky browser usage". Microsoft explains that "workplace stress may lead to uncharacteristic or malicious behavior" by employees, which may "surface as potentially inappropriate behavior" in communication and messaging activities. Visits to "inappropriate or unacceptable" websites may "elevate network security risks", "violate regulatory requirements", expose employers to "legal actions" and "jeopardize current and future business operations and opportunities", according to Microsoft.

In order to identify "insider threats", Microsoft suggests focusing on employees with a "predisposition" or "tendency" to "violate company policies", for example, those who conducted a "policy violation" in the past or, more generally, "disgruntled" employees. To identify "disgruntled" employees, organizations should set up "disgruntlement indicators" or a "dedicated disgruntlement policy". The system provides profiling functionality that promises to detect data leaks by "disgruntled" employees who experience "employment stressors" such as "poor" performance reviews, being placed on a "performance improvement plan" or demotions. Organizations can further investigate suspicious employees and their past behavior, including website visits, file usage and communication content. For "forensic" investigations, employers can access screen recordings and fine-grained user interaction records that show information down to the currently active window at a certain point in time.

While Microsoft's communication monitoring and insider risk systems allow organizations to further investigate employee behavior and communication activity that was assessed as suspicious, the Purview **eDiscovery system** allows them to archive, retrieve and search all "electronic information" associated with certain employees or topics. This includes communication activity and document content processed by systems provided by Microsoft and other vendors. Employers can use the retrieved information for "internal and external investigations" and as "evidence in legal cases" involving "certain executives or other employees". They can target particular employees or search thousands of mailboxes and other contents across the organization. The system helps to identify "persons of interest", discover data "where it lives", find emails, documents and other items "used by people in their day-to-day work tasks" and retrieve data that is "most likely relevant to the case". Mailboxes, chat conversations or file repositories can be put on "legal hold" in order to preserve them and make deletion impossible. eDiscovery investigations can be triggered from the insider risk system when "additional legal review is needed for the user's risk activity".

Microsoft Purview typically processes large amounts of personal data on employees and potentially also on persons outside the organization. It can also analyze extensive personal data on employees to create aggregate reports on suspicious activities. Both the insider risk and the communication monitoring system provide functionality to automatically notify managers about suspicious activities and employees, assign tasks and trigger automated workflows. As briefly discussed in sections 4.5 and 4.6, Purview's **data loss prevention (DLP)** system aims to prevent sensitive information from leaving an organization's IT infrastructure, whether intentionally or not. It monitors communication and file activities across employee devices and other software systems that involve information that was declared sensitive, including in office applications such as Word, Excel and PowerPoint. It raises alerts about suspicious activities focusing on data sharing with external parties. While the DLP system does not create behavioral profiles about employees over time, the insider risk system can share employee risk scores with the DLP system, which can utilize them to single out employees or automatically block them from certain activity.

Microsoft Purview combines a number of software systems for data security, risk profiling and compliance. **Microsoft Sentinel** specifically promises to help organizations detect, investigate and prevent cyberattacks and other cybersecurity threats, including from "malicious insiders". At its core, Sentinel is a "security information and event management" (SIEM) system. As detailed in section 4, it analyzes large amounts of log data from many different

sources "across all users, devices, applications, and infrastructure, both on-premises and in multiple clouds". Sentinel is built on top of other Microsoft systems[380] that can process data on "millions" of activity log records per second, analyze "petabytes" of data and access up to seven years of records. As Sentinel can also process information on suspicious activities and employees from Purview's insider risk and communication monitoring systems, it can be considered a comprehensive security and risk profiling system that can be used for very diverse purposes. Table 7 summarizes how Sentinel processes personal data on employees, according to the findings in section 4.

System, module	Analyzed data categories, data sources	Personal data processing on employees	Purposes, detected behaviors, policies, classifiers
Microsoft Sentinel SIEM and UEBA Analyzes large amounts of log data across many sources to detect cyberattacks, "anomalous" behavior and "insider risks" (section 4.8) *UEBA functionality is also available in "Insider Risk Management" and "Defender for Cloud Apps"* *Can process "Communication Compliance" and "Insider Threat Management" alerts* *Can share alerts with "Insider Risk Management" and other systems.*	• "Millions" of log records per second, logs from up to seven years in the past • Log data about devices, activities, users and employees, e.g. Azure activity, security events, file activity, process/application activity, network activity, office activity, mailbox activity • User account and identity data (e.g. employee name, email address, department, job title, roles, permissions) • Other entity data about devices (hosts), applications (processes), files, IP addresses, URLs, Azure resources, mailboxes and mail messages • "Sensitive data" such as "human resource data" and "geolocation data" • Data from Microsoft systems, e.g. Azure, Microsoft 365, Exchange, SharePoint, Teams, OneDrive, Dynamics 365, Project, GitHub, Purview, Entra, Active Directory, Windows, Windows Firewall, Defender for Endpoint, Identity, IoT and Cloud • Data from third-party systems (e.g. Salesforce, Oracle, SAP, Confluence/Jira, Zoom, Slack, Google Workspace, Snowflake), cloud environments (e.g. Amazon Web Services, Google Cloud, Oracle Cloud) and other sources (e.g. Linux Syslog, SQL databases) • Data from cybersecurity and IT infrastructure systems for network connectivity, virtualization, device/antivirus protection, SIEM, UEBA, threat intelligence and CDN (e.g. Cisco, Juniper, Okta, Citrix, VMware, Symantec, McAfee, Kaspersky, Forcepoint, Exabeam, Fortinet, Rapid7, Proofpoint, SentinelOne, Palo Alto Networks, Digital Guardian, Greynoise Intelligence, Crowdstrike, Akamai, Cloudflare) • Data from Forcepoint's firewall, cloud security, access security and data loss prevention (DLP) systems • Any other data via custom integrations	• Process and analyze extensive personal data on employee behavior • Combine security and activity log data with employee identity data • Monitor and profile employee behavior across the organization over time • Put employees on "watchlists" • Detect "anomalous" or otherwise suspicious behavior based on customizable rules, AI models and "behavioral profiles" across employees, peers and the entire organization • Detect suspicious sequences of seemingly benign activities over time • Continuously calculate risk scores for activities and employees • Single out employees whose behavior was detected as suspicious, rank them by risk and raise alerts • Create tasks and automated workflows from alerts and incidents • Analyze alerts from Microsoft Purview or share data with it • Export alerts or log data • Investigate suspicious employee activity in detail (e.g. files downloaded) • Investigate relationships between entities and identify "similar" incidents • Perform "dragnet" searches for employee activity across log databases • Expand functionality and personal data processing about employees across several software systems • Create aggregate reports on employee activity (e.g. in Microsoft 365, Exchange, SharePoint, Teams)	• Detect cyberattacks, compromised accounts, "external attackers" and "malicious insiders" based on "telemetry and event data" across "users, devices, applications, and infrastructure, both on-premises and in multiple clouds" • Detect security threats, e.g. "password spray", logins from "risky" IP addresses, activity from an Tor IP • Detect "anomalous and therefore suspicious" activities based on user profiles that take "time, peer groups, and expected user activity into consideration" • Detect "non-routine actions", "non-compliant practices", "risky activities" or "suspicious communication" (e.g. "offensive language" via Communication Compliance) • Detect activities such as "anomalous sign-in, "anomalous account creation", "anomalous resource access", "anomalous user activities in Office Exchange", "anomalous user/app activities in Azure", "anomalous web request activity", "unusual network volume anomaly", "anomalous data destruction", "suspicious number of protected documents accessed", "anomalous number of files" downloaded, detect file sharing with "unauthorized" users • Detect "lateral movement" and "data exfiltration" incidents and "advanced multistage attacks" by combining "multiple signals" from various data sources over time • Detect suspicious activities by "user accounts of employees that have high impact value in the organization" or by "terminated employees", i.e. "user accounts of employees that have been, or are about to be, terminated" • Detect suspicious activities by analyzing alerts from Purview Insider Risk Management or Communication Compliance (e.g. detect "offensive language") • Detect suspicious activities in Salesforce event logs, Cisco access logs, Zoom reports, Oracle Cloud event logs, Jira audit logs and many other log data sources • Detect suspicious activities by combining Azure data with data from Forcepoint's security and risk profiling systems to "enrich[es] visibility into user activities" • Detect any kind of suspicious activity based on custom rules and models • Help organizations understand "who the real top risky users are" and whether suspicious users are "compromised, exfiltrating data, or acting as insider threats" by getting "information about who the user is and what is known about them" • Help organizations understand whether a user who was assessed as a risk is an "engineer who often performs unusual activities as part of their job" or a "disgruntled employee who just got passed over for a promotion"

Table 7: Data practices identified in this investigation: Microsoft's Sentinel SIEM and UEBA system

[380] Sentinel is built on top of the "Log Analytics" functionality in Microsoft's "Azure Monitor", which is, in turn, built on top of "Azure Data Explorer".

Microsoft Sentinel can monitor a wide range of activity log records from different sources including "telemetry and event data" from Microsoft systems and data from many other enterprise software systems. It analyzes, for example, security-related log data from Microsoft Azure and Windows. Microsoft's anti-virus system Defender can provide it with access to detailed behavioral data about how employees use files, applications and network connectivity on their computers. Log records processed by Sentinel typically include information that identifies employees and their devices, and, as such, personal data on employees. Specifically, Sentinel can combine security-related log data with information about employee names and job titles from Microsoft's identity and access management systems Entra and Active Directory. More broadly, it can monitor activities in enterprise software systems provided by Microsoft (e.g. Exchange, SharePoint, OneDrive, Teams, Dynamics 365, Project) and many other vendors (e.g. Salesforce, Oracle, SAP, Confluence/Jira, Zoom, Slack). Possible data sources also include other cybersecurity and IT systems that provide network connectivity, firewall functionality, virtualization, anti-virus protection, SIEM and UEBA functionality, threat intelligence or risk profiling (e.g. Cisco, Juniper, Okta, Citrix, VMware, Symantec, McAfee, Kaspersky, Forcepoint, Exabeam, Fortinet, Rapid7, Proofpoint, SentinelOne, Palo Alto Networks, Digital Guardian, Greynoise Intelligence, Crowdstrike, Akamai, Cloudflare). Microsoft emphasizes that organizations can provide Sentinel with "sensitive data" such as "human resource" information or "geolocation data".

Based on all this activity log data and digital profiling, Sentinel aims to help organizations detect and prevent cyberattacks, compromised accounts and other security threats. It includes UEBA functionality that creates "behavioral profiles" in order to identify "anomalous and therefore suspicious" activities associated with users, employees and other "entities" such as devices, files and applications. Sentinel's UEBA module continuously calculates risk scores for activities, users and employees. It singles out those who are assessed as suspicious, ranks them by risk and raises alerts. The calculated risk score represents the "sum of all the user's risky activities over the last week", for example, and indicates "how risky a user is relative to other users" in the organization. To identify unusual activities, Sentinel analyzes behavioral data across the organization over time. It promises to detect, for example, "anomalous" logins, "anomalous" access to documents or file repositories, "anomalous user activities" in Exchange, "anomalous user/app activities" in Azure, "anomalous" web activities, "anomalous" file downloads, "anomalous" file deletions and "anomalous" file sharing. To identify sophisticated threats, it can analyze sequences of activities over time. The system can detect suspicious activities in Salesforce event logs, Cisco access logs, Zoom reports, Oracle Cloud event logs, Jira audit logs and many other log data sources. To detect cybersecurity threats, it can also utilize alerts about "suspicious communication" such as the use of "offensive language" by employees.

More broadly, Sentinel promises to detect "non-routine actions", "non-compliant practices" and other "risky activities" that indicate a threat to the organization, including by employees who are considered potential "insider threats". Organizations can focus on certain groups of employees by putting them on a "watchlist". The system provides, for example, pre-configured watchlists for "VIP users" that have a "high impact value in the organization" or "terminated employees" that "have been, or are about to be, terminated". Based on alerts about suspicious activities or employees, it can also create tasks and trigger automated workflows. After Sentinel has detected suspicious activities, organizations can further investigate the corresponding user accounts and employees, their past behavior, relationships with other user accounts and "similar" incidents. They can "get information about who the user is and what is known about them" and "investigate all related activities" in order to understand whether user accounts are "compromised" or "acting as insider threats". Microsoft emphasizes that Sentinel helps organizations understand whether a suspicious user is an "engineer who often performs unusual activities as part of their job" or a "disgruntled employee who just got passed over for a promotion". In addition, Sentinel provides extensive functionality to perform "dragnet" searches across large amounts of activity log data according to custom criteria.

Sentinel can import data from several other cybersecurity and risk profiling systems provided by Microsoft, such as Defender for Identity, Defender for IoT, Defender for Cloud and Purview. This includes alerts about suspicious employees from Purview's insider risk and communication monitoring systems. It can also import data about suspicious activities from Forcepoint's security and risk profiling systems, which can enrich "visibility into user activities", enable "further correlation with data from Azure" and improve "monitoring" inside Sentinel, according to Microsoft. In turn, Microsoft's UEBA functionality can be utilized in Purview's insider risk system and Microsoft's "Defender for Cloud Apps" product.

6.3 IBM and Teramind

This case study briefly examines cybersecurity and risk profiling software provided by other vendors, as detailed in section 5. **IBM's "QRadar"** system offers UEBA, SIEM and insider threat detection functionality similar to Forcepoint and Microsoft. It can monitor "millions" of activity records from event logs, employee devices and network data across 450 data sources and 370 applications and analyze them "against historical data to uncover known and unknown threats", from cyberattacks to "insider threats" to non-compliance with "enterprise, industry, and regulatory" policies. It creates "risk profiles" and "risk scores" about users and employees in order to detect "anomalous behavior". Monitored behaviors can include visits to "gambling" or "lifestyle" websites and "negative sentiment" in employee communication, including in email contents. QRadar also provides functionality for "forensic" investigations and searches across large amounts of activity log data. In addition, IBM offers a **communication monitoring** system for financial organizations, which scans employee communication via email, chat and voice calls to prevent non-compliant behavior, misconduct and financial crimes. It aims to identify communication activities that indicate insider trading, confidentiality violations or efforts to "recruit co-conspirators". It also promises to detect "emotions" and emotion "anomalies" such as "anger" and "sadness" in conversations.

Teramind, a smaller US-based software vendor, offers a far-reaching employee surveillance system that combines cybersecurity, insider threat detection and "compliance management" with productivity monitoring. Based on software installed on the computers and devices of employees, the system monitors application usage, website visits, email and message communication, online meetings, file usage, keyboard activity and screen activity. Teramind calculates risk scores for employees, categorizes the monitored activities as either "productive" or "unproductive" and provides a wide range of employee rankings. Employers can use it to "capture, analyze and control user desktop activity for any use case", according to Teramind. They can access screen recordings, search for employees who entered certain keystrokes and identify "unproductive employee activity" and "low-performing employees".

7. Discussion and concluding remarks

This case study shows that today's cybersecurity and risk profiling systems provide far-reaching surveillance capabilities. As data collection in the workplace has become ubiquitous, these systems can utilize an increasing amount of data on employees that is often originally processed for other purposes. They can monitor how employees use files and applications, the websites they visit, their searches, email and chat conversations, voice calls, video meetings, how they physically access buildings and rooms, their geolocation, performance reviews and even clipboard, keyboard and screen activity. Organizations can potentially use these systems to monitor **everything employees do or say**, profile their behavior and single them out for further investigation.[381]

The **cybersecurity, insider risk and communication monitoring systems** examined in this study[382] process extensive personal data on employee behavior and communication from many sources. They serve different purposes ranging from the prevention of cyberattacks and the identification of employees who are considered a threat to the organization to the detection of misconduct and otherwise undesirable behavior. As these systems typically share data with each other, they become combined security and risk surveillance systems, which not only promise to *detect* incidents but to *prevent them before they occur*. Similar to **predictive policing** technology (Kuldova, 2022), they offer to detect "anomalies" and otherwise suspicious behaviors that could indicate future incidents. Because of the scale and depth of data collection and their rich capabilities for behavioral profiling over time they can be considered **corporate mass surveillance** or **dragnet surveillance** systems. In addition to detecting suspicious behavior, they help organizations automate how to respond to alerts, from notifying managers to automatically blocking employees from performing certain activities. They help them investigate past employee activity and allow them to search for certain kinds of behavior across large amounts of activity data, both in real time and over long periods of time.[383]

The systems examined in this case study continuously calculate risk scores for employees and provide profiling mechanisms that make **intrusive inferences** about employees. Forcepoint offers to assess whether employees are in financial distress, show "decreased productivity" or intend to leave the job, how they communicate with colleagues and whether they access "obscene" content or exhibit "negative sentiment" in their conversations.[384] Microsoft's communication monitoring system offers to detect everything from "profanity", "inappropriate" language, harassment and discrimination to corporate sabotage, data leaks, money laundering, bribery, conflicts of interest and "workplace collusion". Microsoft's insider risk system additionally promises to identify "disgruntled employees" and suspicious activities such as visiting "inappropriate or unacceptable" websites. The company's Sentinel cybersecurity system helps organizations understand whether a user whose behavior was identified as suspicious is a "disgruntled employee who just got passed over for a promotion".[385] In addition to the systems' built-in data sources and profiling mechanisms, organizations can create their own risk policies that can **detect any type of behavior or communication pattern** based on custom data sources, rules, keywords and AI-based classifiers.[386] It

[381] See summary in section 6
[382] Note: The research in this case study refers to products offered by Forcepoint up until late 2023 (see section 3)
[383] Ibid.
[384] See section 3.2
[385] See sections 4.1, 4.2 and 4.8
[386] See summary in section 6

is quite easy to set up a policy or "watchlist" that aims to detect, for example, workplace organizing or employees who complain about working conditions.

Of course, organizations are required to prevent cyberattacks and protect information security, which represents a legitimate purpose for data collection and analysis. To a varying extent, this also applies to measures against employee misconduct, from violations of regulatory and organizational policies to criminal conduct. Nevertheless, **the findings of this case study raise serious concerns** about potential misuse by employers, intrusive surveillance, disproportionate data processing, inaccurate risk assessments, arbitrary suspicion and the use of information about employee performance for risk profiling:

- Employers can potentially misuse the powerful surveillance capabilities provided by these systems to **unethically or illegally spy on employees** or even on employee representatives, organized labor and works councils. They can use them to silence internal dissent or to find evidence of inappropriate behavior that would justify terminations or arbitrary disciplinary action. The notion of employees as potential "insider threats" is ambiguous. Intel refers to "activists" who are "highly motivated supporters of a cause" as potential cybersecurity threats.[387] Forcepoint suggests targeting "disgruntled employees" who have had a "huge fight with the boss" and "internal activists" who leak information to journalists.[388] Microsoft suggests focusing on employees with a "predisposition" to "violate company policies" and provides functionality to target "disgruntled employees".[389] More broadly, employers can potentially utilize these systems for **excessive behavioral policing** and control.
- Even if organizations do not intentionally misuse these systems, employees are **put under general suspicion** of being "risks" and "threats" or even seen as guilty until proven innocent (Kuldova, 2022). Constant surveillance of behaviors and conversations undermines **employee privacy**, which can affect how those who are the subjects of surveillance act and communicate. It potentially undermines employees' **autonomy, liberty, freedom of expression and human dignity**. The surveilled may act more cautiously and avoid attention, information seeking and opinion sharing (Solove, 2006; Raab, 2014; Büchi et al., 2022). Employee surveillance can **sow distrust**. Even if well-intentioned, it may "alienate the very people whose commitment and trust" organizations "are trying to secure" (Chory et al., 2015). Surveillance generally increases the **power and information asymmetry** between organizations and employees (Andrejevic, 2014).
- The security and risk profiling systems examined in this study raise concerns about **overly intrusive and disproportionate** processing of personal data on employees in relation to the purposes of processing. This generally concerns the scale, scope, depth and frequency of data analysis across different sources.[390] For many employees, monitoring communication, file, application and browser activity is equivalent to monitoring almost all work activities. Monitoring keyboard and screen activity is even more intrusive. Analyzing sensitive information such as the contents of communication raises specific concerns. Depending on the purpose, continuous behavioral profiling can generally be considered overly intrusive and disproportionate. This may be even more the case if the profiling results in **sensitive inferences or assessments**, for example, with regard to emotions (e.g. "negative" sentiment, anger, sadness)[391] and similar characteristics (e.g. disgruntledness),[392]

[387] See section 2.2
[388] Ibid.
[389] See section 4.1
[390] See summary in section 6
[391] See sections 3.2 (Forcepoint), 4.2.2 (Microsoft) and 5 (IBM)
[392] See section 4.1 (Microsoft)

personal characteristics (e.g. personal finances, intent to leave the job)[393] or sensitive classifications of website visits (e.g. health, drugs, abortion, political organizations).[394] Forcepoint's data sources can also include external information about "criminal history" and "financial distress". It can analyze "off hour" and "weekend" activity, and as such, it potentially monitors the private lives of employees.[395]

- Singling out employees as risks that deserve further investigation based on extensive behavioral data and AI-based profiling raises concerns about **inaccurate risk assessments and arbitrary suspicions**. Risk scoring algorithms that predict potentially malicious behavior are typically opaque and unreliable by design (see e.g. boyd and Crawford, 2012; Citron and Pasquale, 2014). Like other predictive policing technologies, security and risk profiling systems used by employers may create "inaccurate, skewed, or systematically biased data" and "flawed" predictions (Richardson et al., 2019). Automated responses to alerts may aggravate the possible harms. When employees or certain groups of employees get inaccurately accused of "anomalous" or otherwise suspicious behavior by an organization's cybersecurity, compliance or human resource departments or when they get automatically blocked from certain activity, this may lead to **Kafkaesque experiences**.
- Microsoft acknowledges that its cybersecurity and risk profiling systems may create "false positives",[396] i.e. inaccurate alerts about employees and their behavior, which is why it provides a wide range of functionality to prioritize, review and investigate alerts. Reviewers, analysts, investigators or managers may, however, still treat employees who are often accused of suspicious behavior by the system differently than others, even if these accusations regularly turn out to be inaccurate. Microsoft also acknowledges that its Sentinel system may create large amounts of records on behavioral anomalies, which makes the data "notoriously very noisy".[397] Microsoft's communication monitoring system offers the option to quickly train custom classifiers by feeding it with 50 text samples that "represent the type of content" an organization wants to detect.[398] It is doubtful whether this can lead to reliable results. The reliability of AI-based emotion detection is generally questionable (Stark and Hoey, 2021; Corvite et al., 2023). Taken together, the examined systems raise questions about effectiveness, accuracy and reliability at several levels.
- Several systems analyze data about **employee performance and productivity** in order to make security and risk assessments. Both Forcepoint and Microsoft prominently suggest utilizing HR information from performance reviews.[399] Employees whose performance was assessed as "poor" or otherwise inadequate are considered particularly suspicious and receive extra scrutiny. As such, security and risk surveillance may indirectly allow organizations to implement more rigid forms of performance control, which can have the effect of turning employees into "disgruntled" employees who pose a risk to the organization. The inclusion of HR information about employees who were demoted or put on performance improvement plans has similar implications. Forcepoint goes one step further by directly utilizing productivity monitoring for risk assessments.[400] In addition, Forcepoint's security and risk profiling systems can display information about the amount of time suspicious employees spent on activities such as using programs, moving files, visiting websites, searching the

[393] See section 3.2 (Forcepoint)
[394] See section 3.4 (Forcepoint)
[395] Ibid.
[396] See e.g.: https://learn.microsoft.com/en-us/purview/insider-risk-management-settings-intelligent-detections, https://learn.microsoft.com/en-us/purview/insider-risk-management-settings-detection-groups, https://learn.microsoft.com/en-us/purview/communication-compliance-solution-overview, https://learn.microsoft.com/en-us/azure/sentinel/false-positives [20.5.2024]
[397] See section 4.8.2
[398] See section 4.2.1
[399] See summary in section 6
[400] Forcepoint's "decreased productivity" risk model promises to identify employees who spend "a large amount of time doing non work related tasks" (see section 3.2)

web and using email.[401]

- Several cybersecurity firms are closely affiliated with the defense and intelligence sector. Forcepoint was until recently owned by the defense giant Raytheon. Its UEBA technology that promises to detect "anomalous" behaviors originates from a company that received funding from the CIA's venture capital firm In-Q-Tel and was acquired by Forcepoint in 2017.[402] Exposing employees to **military-grade surveillance technology** raises concerns, even more so as Forcepoint appears to process large amounts of data across organizations globally. Forcepoint claims that its insider risk system monitors more than one million devices. Across its cybersecurity solutions, it claims to analyze five billion activity records per day from 900 million devices.[403]

The security and risk profiling systems examined in this case study provide intrusive surveillance for **very different purposes**. The technology is becoming broadly available for organizations in many sectors. Personal data processing that is legitimate for some purposes and employees may be disproportionate for other purposes and employees:

- The finding of this study suggests that the **boundaries** between cybersecurity, fraud and theft prevention, the protection of customer data and trade secrets and the enforcement of "compliance" with laws, guidelines, policies, codes of conduct and other organizational rules are becoming increasingly blurred. Data from systems for cybersecurity, access control and device management is being used to detect employee misconduct, whether intentional, unintentional, negligent or otherwise undesirable. Anti-virus software that blocks access to malware-infected websites is also used to detect undesirable web activity. Behavioral profiles on web activity, "offensive language" in communication, and HR information such as performance reviews contribute to assessments of employees as potential "insider threats". The examined security and risk profiling systems provide intrusive **behavioral surveillance across very different purposes** ranging from the detection of inappropriate language in conversations to the prevention of highly consequential threats such as cyberattacks and criminal conduct.[404] As the same systems can easily be used for very different purposes, this bears the risk of disproportionate personal data processing being employed for purposes that do not justify excessive behavioral surveillance.
- Several systems examined in this study offer specific risk profiling functionality for organizations in the **financial sector** that addresses insider trading, money laundering and other forms of misconduct that is subject to government regulations.[405] Forcepoint promises to detect attempts to deceive polygraph tests or evade security clearance checks, which clearly points to the company's close affiliation with the **defense and intelligence sector**. Customers of Forcepoint's security and risk profiling software extend, however, beyond defense and banking — they include organizations in many sectors, including some located in Europe.[406] Forcepoint provides intrusive profiling mechanisms that are not built into other systems. Microsoft's security and risk profiling systems, however, offer almost the same functionality and are **easily available to many organizations that already use Microsoft software**. While Microsoft does not provide built-in risk profiling mechanisms related to personal finances and stopped rolling out its "leavers" classifier after a public outcry in 2022,

[401] See section 3.1
[402] See section 2.4
[403] See section 3.6
[404] See summary in section 6
[405] See sections 3.6 (Forcepoint), 4.2.1 (Microsoft) and 5 (IBM)
[406] See section 3.6

organizations can easily implement custom detections for almost any type of behavior.[407] Microsoft's communication monitoring system already offers a wide range of built-in detections for many purposes ranging from "inappropriate" language to bribery, conflicts of interest and corporate sabotage.[408] Consequently, intrusive cybersecurity and risk profiling is now broadly available for both large and smaller organizations across all sectors.

- The broad availability of these technologies raises concerns about organizations **disproportionally deploying them across their entire staff**. Not all employees are equal. It may be justified to subject some high-salaried employees to more intrusive surveillance, for example, a server administrator whose user account is vulnerable to consequential cyberattacks or a higher-level executive who is vulnerable to coercion. Applying intrusive surveillance to some employees with access to specifically sensitive resources certainly does not automatically justify applying the same level of surveillance to large groups of employees or an organization's entire staff. Even if the monitoring is not targeted at all employees, the UEBA technology provided by Forcepoint and Microsoft may process extensive personal data on employees across departments or an entire organization in order to build behavioral profiles that allow for the detection of "anomalous" behavior.[409] Microsoft also offers to process data on behaviors and communication activities across the organization for aggregate reports.[410]
- As cybersecurity threats are considered existential and potentially catastrophic, they justify intrusive surveillance that would otherwise not be accepted (Da Silva, 2022). The expansion of the scope of risk surveillance from cybersecurity to "insider threats" to compliance with all kinds of organizational policies can be considered a typical example of **function creep** (Koops, 2021). As soon as behavioral surveillance for cybersecurity is implemented, it is only a small step to applying the same technology for other purposes. The findings of this study suggest that today's cybersecurity and risk profiling systems generally contribute to the expansion and **normalization of pervasive employee surveillance**.
- Microsoft recommends that organizations monitor the communication activities of "all" employees at least for "harassment or discrimination detection".[411] This is problematic for three reasons. First, communication monitoring is intrusive. Second, implementing it for one purpose opens the door to implementing it for other purposes. Third, it is doubtful whether monitoring all employee communication is actually an appropriate solution to harassment and discrimination — it may rather represent an intrusive **technological pseudo-fix** for issues that are deeply embedded in corporate cultures and deserve serious attention at all levels of an organization.
- As briefly examined in section 5, the employee surveillance system provided by Teramind openly combines risk and productivity monitoring. Because of its intrusive multi-purpose functionality and the lack of adequate safeguards, it is unlikely that European employers can lawfully use it. Nevertheless, Teramind appears to sell to Europe, and it is particularly concerning that its system can be easily used by even the smallest businesses.

Organizations can implement and use the cybersecurity and risk profiling systems examined in this case study in very different ways. They can customize data sources and analysis functionality, decide to apply behavioral monitoring to some or many employees and implement **more or less effective safeguards** that promise to address employee privacy, data protection, misuse and other concerns:

[407] See section 4.2.1
[408] Ibid.
[409] See summary in section 6
[410] Ibid.
[411] See section 4.2

- The risk profiling systems provided by Forcepoint and Microsoft offer to replace employee names with **pseudonyms** in the user interface. This can prevent unnecessary access to directly identifiable personal data for those who initially review alerts about suspicious employees. As employee identities can still be accessed if required,[412] it does not fundamentally change anything in the systems' capability to single out employees based on extensive behavioral profiling. Microsoft's Sentinel cybersecurity system does not offer pseudonymization in the user interface.[413] While organizations can abstain from utilizing information about employee identity, the Sentinel software documentation prominently describes use cases that require it.[414]
- Both Forcepoint and Microsoft provide a number of **permissions and roles** that determine who can access what kind of data and who can modify what is being monitored. While Microsoft provides a wide range of pre-configured roles, it undermines its own security and privacy promises by suggesting, for example, that using a single role for all insider risk functionality would be the "easiest way to quickly get started" and can be a "good fit for organizations that don't need separate permissions".[415] When sensitive data about suspicious employees, as determined by Microsoft's insider risk or communication monitoring systems, is shared with the company's Sentinel cybersecurity system,[416] it leaves the scope of some safeguards built into the former systems. Most systems that were examined make it possible to export almost any data,[417] which is where any available safeguards cease to apply.
- As detailed in section 4.9, Microsoft's **audit log** can act both as an additional employee surveillance tool and as a means to help organizations or worker representatives monitor and prevent potential misuse of its security and risk profiling systems for inappropriate employee surveillance.

As outlined in section 4.10, Microsoft makes a variety of data protection claims. Forcepoint stated in 2017 that its risk profiling systems would not put employees under general suspicion and that organizations could comply with all requirements of a German works council.[418] A comprehensive assessment of these measures regarding the GDPR and labor law in Germany and other European countries is beyond the scope of this study. Nonetheless, Forcepoint and Microsoft offer intrusive functionality out of the box and sometimes recommend organizations implement the more intrusive options or even incentivize them to expand employee surveillance:

- Both Forcepoint and Microsoft offer **built-in profiling mechanisms that can be considered intrusive**. Microsoft's communication monitoring system provides a wide range of intrusive classifiers for very different purposes.[419] Its insider risk system can utilize information about the "use of offensive language", "risky browser usage" and "disgruntled" employees.[420] Forcepoint recommends setting up a number of highly intrusive behavioral profiling models that promise to assess whether employees are in financial distress, show "decreased productivity" or intend to leave the job, how they communicate with colleagues, and whether they access "obscene" content or exhibit "negative sentiment" in their conversations.[421]

[412] See sections 3.7 (Forcepoint) and 4.10 (Microsoft)
[413] See section 4.10
[414] See section 4.8
[415] See sections 3.7 (Forcepoint) and 4.10 (Microsoft)
[416] See section 4.3
[417] See summary in section 6
[418] See section 3.7
[419] See section 4.2.1
[420] See section 4.1
[421] See section 3.2

- Both vendors suggest that companies **deploy intrusive surveillance across the entire organization**. Forcepoint recommends that monitoring "should include all employees at a company".[422] Microsoft recommends monitoring "all" employees in an organization at least for "harassment or discrimination detection",[423] which opens the door for comprehensive communication monitoring for other purposes.
- Microsoft **systematically incentivizes organizations** to implement far-reaching employee surveillance by offering them the ability to quickly analyze massive amounts of personal data on employee behavior and communication and awarding them "points" which promise to measure their "progress towards completing recommended actions". For example, Microsoft awards organizations these points for setting up the insider risk and communication monitoring systems. Subsequently, it provides them with recommendations that, for example, suggest to "start monitoring communications" to detect "inappropriate" content "now".[424]

In conclusion, the cybersecurity and risk profiling systems offered by Forcepoint and Microsoft are designed to utilize almost all available personal data on employee behavior and communication across different purposes. Their software documentation and promotional materials make many suggestions about intrusive monitoring and profiling. As the previous section shows, they make a wide range of suggestions about how to use these systems, from detecting "anomalous" and "inappropriate" employee behavior to considering "disgruntled employees" and those with a "predisposition" to "violate company policies" as a major risk that justifies far-reaching surveillance.[425] Of course, the employers who actually deploy these systems are mainly responsible for how they implement them. Software vendors, however, shape how they can be used and thus how employees are affected. The enterprise software giant Microsoft carries a specific responsibility here. Its software is easily available to many organizations who already use Microsoft 365 or other systems provided by the company. The findings of this case study suggest that Microsoft promotes extensive employee surveillance as a technological solution to many problems. The idea that technology and data can fix everything has often been criticized as **tech or data solutionism** (Morozov, 2013). Unfortunately, this is exactly how most software for cybersecurity and risk profiling is advertised.

Employee representatives and work councils can only be advised to carefully discuss and negotiate the potential deployment of a SIEM, UEBA, DLP, insider risk or communication monitoring system with the employer. What kinds of personal data and behavioral profiling are really necessary and appropriate for which purposes and employee groups? Who has access to what kinds of reports and investigation functionality? How to prevent misuse for other purposes and which safeguards, checks and balances can be implemented? Extremely intrusive functionality such as monitoring keyboard and screen activity should be disabled or only used under very strict circumstances for employees who deal with highly sensitive resources. Any monitoring relating to the private lives of employees such as their financial situation and any unproven technology such as emotion detection should be disabled. In many cases, it can make sense to use completely separated systems for separate purposes and avoid vendors that are either closely affiliated with the defense industry or provide very comprehensive cross-purpose systems.

[422] Ibid.
[423] See section 4.2
[424] See sections 4.1.1 and 4.2
[425] See section 4.1

List of figures

References

Andrejevic, M. B. (2014): The big data divide. International Journal of Communication, 8(1), 1673–1689

Alneyadi, Sultan; Sithirasenan, Elankayer; Muthukkumarasamy, Vallipuram (2016): A survey on data leakage prevention systems. Journal of Network and Computer Applications, Volume 62, 2016, Pages 137-152, DOI: 10.1016/j.jnca.2016.01.008

Cardoso, N. A. D. M. L. (2021): User behavior analytics in the contact center: Insider threat assessment and fraud detection. Dissertation, July 2021. Online: https://estudogeral.uc.pt/handle/10316/96092

Chory, Rebecca; Vela, Lori; Avtgis, Theodore (2015): Organizational Surveillance of Computer-Mediated Workplace Communication: Employee Privacy Concerns and Responses. Employee Responsibilities and Rights Journal. 28. DOI: 10.1007/s10672-015-9267-4.

Christl, Wolfie (2021): Digitale Überwachung und Kontrolle am Arbeitsplatz. Von der Ausweitung betrieblicher Datenerfassung zum algorithmischen Management? Eine Studie von Cracked Labs, 2021. Available online: https://crackedlabs.org/dl/CrackedLabs_Christl_UeberwachungKontrolleArbeitsplatz.pdf

Citron, Danielle K. & Pasquale, Frank (2014): Essay, The Scored Society: Due Process for Automated Predictions, 89 Wash. L. Rev. 1 (2014)

Corvite, Shanley; Roemmich, Kat; Rosenberg, Tillie Ilana; Andalibi, Nazanin (2023): Data Subjects' Perspectives on Emotion Artificial Intelligence Use in the Workplace: A Relational Ethics Lens. Proc. ACM Hum.-Comput. Interact. 7, CSCW1, Article 124 (April 2023), DOI: 10.1145/3579600

boyd, danah, & Crawford, K. (2012): Critical questions for Big Data: Provocations for a cultural, technological, and scholarly phenomenon. Information, Communication & Society, 15(5), 662–679. DOI: 10.1080/1369118X.2012.678878

Büchi, M.; Festic, N.; Latzer, M. (2022): The Chilling Effects of Digital Dataveillance: A Theoretical Model and an Empirical Research Agenda. Big Data & Society, 9(1). DOI: 10.1177/20539517211065368

Gartner (2015): Magic Quadrant for E-Discovery Software. Gartner, 18 May 2015

Gartner (2018): Market Guide for User and Entity Behavior Analytics. Gartner, 23 April 2018

Gartner (2020): Market Guide for Insider Risk Management Solutions. Gartner, 29 December 2020

Gartner (2022): Magic Quadrant for Security Information and Event Management. Gartner, 10 October 2022

Gartner (2023): Market Guide for Digital Communications Governance. Gartner, 13 November 2023

Gelles, Michael G. (2016): Insider Threat: Prevention, Detection, Mitigation, and Deterrence. Elsevier Science, 2016

González-Granadillo, G.; González-Zarzosa, S.; Diaz, R. (2021): Security Information and Event Management (SIEM): Analysis, Trends, and Usage in Critical Infrastructures. Sensors 2021, 21, 4759, DOI: 10.3390/s21144759

Khaliq, S., Abideen Tariq, Z. U. and Masood, A. (2020): Role of User and Entity Behavior Analytics in Detecting Insider Attacks. 2020 International Conference on Cyber Warfare and Security (ICCWS), Islamabad, Pakistan, 2020, pp. 1-6, DOI: 10.1109/ICCWS48432.2020.9292394

Koops, Bert-Jaap (2021): The concept of function creep. Law, Innovation and Technology, 13:1, 29-56, DOI:

10.1080/17579961.2021.1898299

Kovacs, Laszlo (2018): National Cyber Security as the Cornerstone of National Security. Land Forces Academy Review. 23. 113-120. DOI: 10.2478/raft-2018-0013

Kuldova, Tereza Østbø (2022): Compliance-Industrial Complex. The Operating System of a Pre-Crime Society. Palgrave Macmillan, 2022

Kuldova, T. Ø. and Nordrik, B. (2023): Workplace investigations, the epistemic power of managerialism and the hollowing out of the Norwegian model of co-determination. Capital & Class, 2023, 1-26. DOI: 10.1177/03098168231179971, available online: https://oda.oslomet.no/oda-xmlui/handle/11250/3072224

Morozov, Evgeny (2013). To Save Everything, Click Here. New York, Public Affairs, 2013.

Raab, Charles (2014): Surveillance: Effects on Privacy, Autonomy and Dignity. In D. Wright, & R. Kreissl (Eds.), Surveillance in Europe (pp. 259-268). Routledge.

Richardson, Rashida; Schultz, Jason; Crawford, Kate (2019): Dirty Data, Bad Predictions: How Civil Rights Violations Impact Police Data, Predictive Policing Systems, and Justice. N.Y.U. Law Review Online, 192.

Riesenecker-Caba, Thomas (2023): Datenmonster (Home-)Office Arbeitsplatz? Im Spannungsfeld von IT- und Datensicherheit und versteckter Kontrolle. Forba, Jänner 2023. Available online: https://www.forba.at/wp-content/uploads/2023/01/Datenmonster-Home-Office_IT-und-DatensicherheitFORBA-Jan2023-1.pdf

Sachowski, Jason (2018): Digital Forensics and Investigations: People, Process, and Technologies to Defend the Enterprise. CRC Press, 2018

Salitin, M. A. and Zolait, A. H. (2018): The role of User Entity Behavior Analytics to detect network attacks in real time. 2018 International Conference on Innovation and Intelligence for Informatics, Computing, and Technologies (3ICT), Sakhier, Bahrain, 2018, pp. 1-5, DOI: 10.1109/3ICT.2018.8855782

Solove, D. J. (2006): A Taxonomy of Privacy. University of Pennsylvania Law Review, 154(3), 477–564. DOI: 10.2307/40041279

Stark, Luke and Jesse Hoey (2021): The Ethics of Emotion in Artificial Intelligence Systems. In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21). Association for Computing Machinery, New York, NY, USA, 782–793, DOI: 10.1145/3442188.3445939

Vacca, John R. (2017) (eds): Computer and Information Security Handbook. Third Edition. Morgan Kaufmann, 2017