eccenca is a German software company specializing in knowledge graph based solutions with more than a decade of experience. Our product eccenca Corporate Memory helps enterprises break down data silos and fuse their data assets with expert knowledge to ultimately enable business owners to make decisions based on a holistic view. eccenca partners with Ontotext and uses its GraphDB as the engine for its knowledge platform.

Following, we'd like to introduce you to our product and its use within the Cybersecurity space.

User Journey and Functional Overview

We aim to make building and using our knowledge platform as easy and intuitive as possible. The following picture provides a functional overview of eccenca Corporate Memory. In the diagram below you see our platform's four main categories of capabilities.

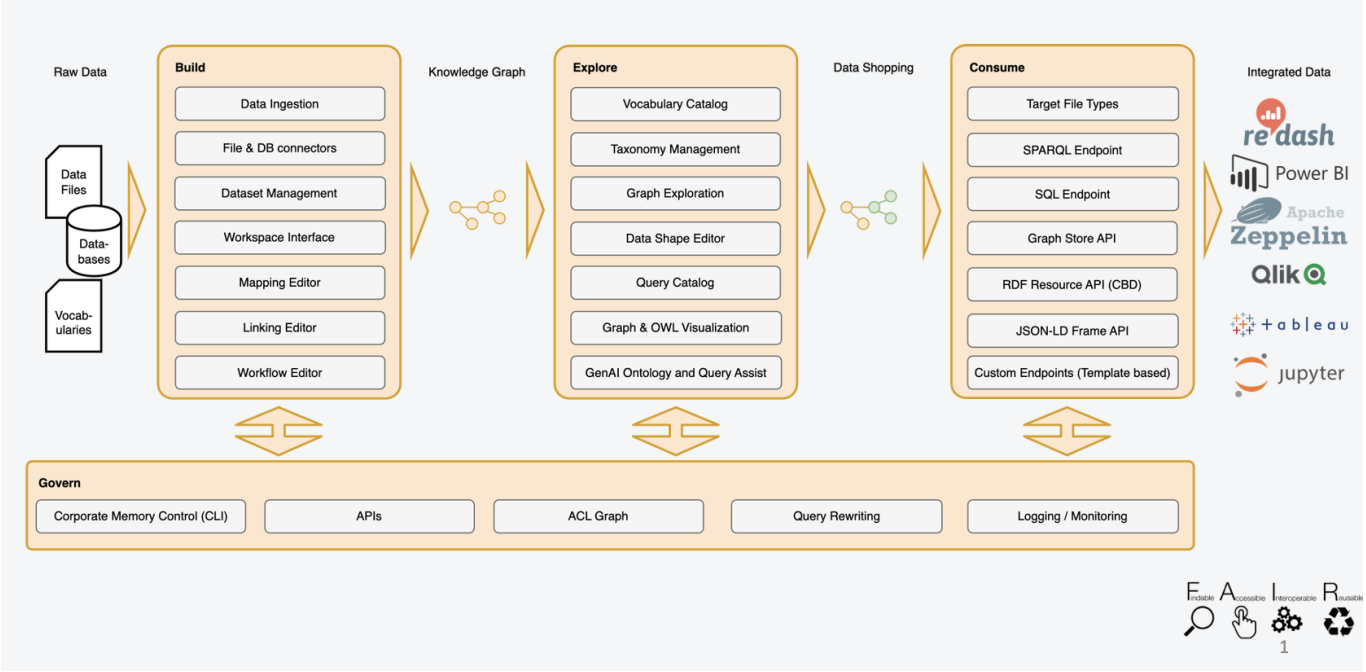

"Build" is where we typically start our projects. We connect to data sources, take vocabularies, and build workflows that turn your non-graph datasets into a knowledge graph. When you have a knowledge graph in place, the second category "**Explore**" offers various

Copyright: Ontotext New York | Sofia | Basel
Reach out to us: info@ontotext.com

features and functionalities for interacting with it; while "**Consume"** provides different APIs and integration endpoints to anchor this application in the existing IT infrastructure. Finally, there's the category of "**Govern"**, which is a shared backend functionality layer that we use – among other functions – to provide a way of doing access control of your graph-based solution.

A recent GenAI-driven capability we've added is our Ontology Assistant. Very often, ontology work is the first milestone a project needs to master. This Assistant offers a chat-like interaction, where you can describe the business objects in your domain and it helps you to overcome the problem of the blank sheet. As mentioned before, building a knowledge graph and making use of insights should be convenient for everyone interacting with our product.

Another new feature is our GenAI-driven Query Generation. Again, it's a chat-like interface that translates your plain English request for data into a technical SPARQL query that can be executed directly and results can be taken from the graph.

Cybersecurity and Vulnerability Management Solution

Our knowledge graph based solutions can be applied in many use cases but here we want to talk about Cybersecurity. We see Cybersecurity as an important market segment in which we have been developing neuro-symbolic AI solutions together with leading customers from industry and government for some time. For instance, we are partnering with Airbus in this field. Based on our experience, we have found three main elements in IT security where we can be of help:

- Prevention of attacks (Firewalls, Packet Inspection, etc.)
- Threat detection and mitigation
- Hardening of existing software

What's typical of the Cybersecurity domain is that it is very knowledge-rich but doesn't care much about connecting and contextualizing all this information. It also doesn't focus on reuse, so many activities of tracing down attacks are one-off events. Rarely are they well documented and well anchored in other knowledge assets of this domain. So, this is an area where we can bring significant benefits to the table.

Copyright: Ontotext New York | Sofia | Basel
Reach out to us: info@ontotext.com

Vulnerability Management Solution

Today, we want to focus on the third category of solutions – hardening of existing software and, more specifically, Common Vulnerabilities and Exposure (CVE).

What we offer is a free add-on package that can be used in combination with an eccenca.my instance for the specific use case of managing the CVEs. We consider it as the nucleus of a potentially growing and nurturing family and ecosystem of Cybersecurity solutions and components. In the future, this can also extend to other areas like threat detection, ticketing about threats, threat intelligence, and threat forensics, where this whole nucleus can play an important role.

As we said, the Cybersecurity domain is data-heavy, but not models-heavy, so it can benefit from a good data strategy and a semantic data approach. Having a better paradigm to manage different signals is a big time-saver and can make the life of a Cybersecurity analyst much easier.

The Solution: In a Nutshell

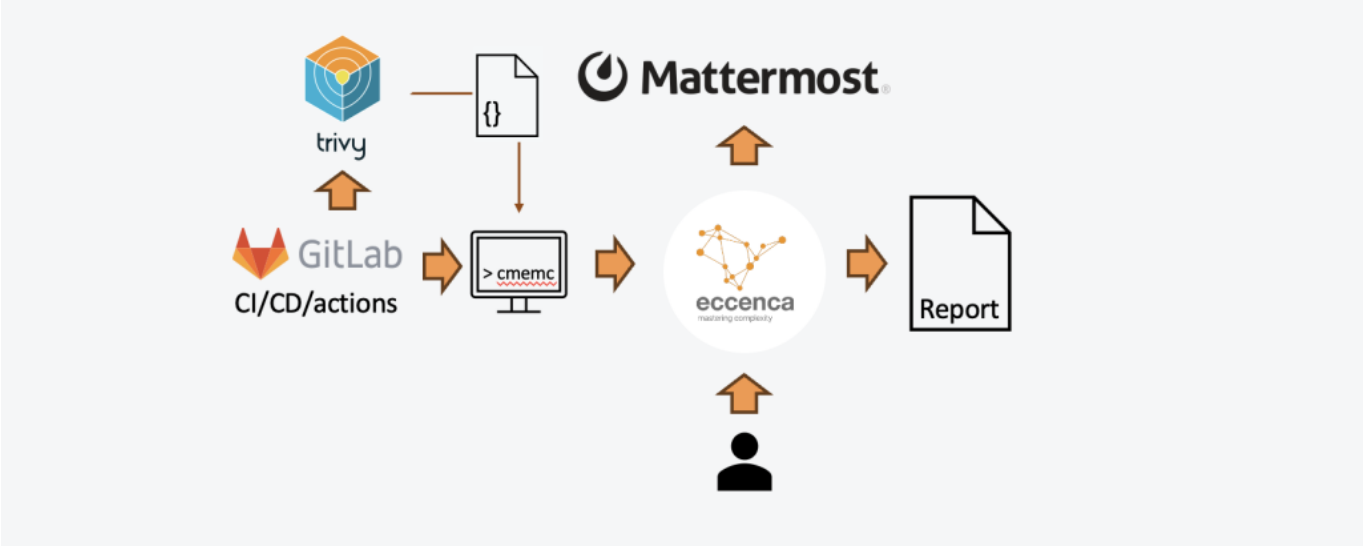

As you can see in the diagram above, the solution is centered around eccenca Corporate

Copyright: Ontotext New York | Sofia | Basel
Reach out to us: info@ontotext.com

Memory as the integral brain. It starts by scanning software packages – a process driven by two third-party elements. It runs as a CI/CD pipeline in GitLab that can be triggered daily or on demand. It invokes a Trivy scanner, which takes for inspection a list of software artifacts (namely, Docker images). The Trivy scan results (JSON documents) are then taken and fed by cmemc (our command line automation utility) and sent to a Corporate Memory workflow, which fills our vulnerabilities knowledge graph with additional elements.

A specific workspace is configured in eccenca Corporate Memory and is the entry point for our software engineers. The pipeline connects to a Mattermost (our internal communication channel) chatbot that informs the developers' community about newly found CVEs. It asks our experts to outline potential mitigation strategies and evaluate the risk vector of these newly found CVE.

As a result of this annotation process, the next time this pipeline is run, an updated report is generated. It highlights the strategies our experts outlined for the CVEs that were found.

The Solution: Walk the Talk

The figure below shows the internal data model shown in our built-in visual graph viewer. You can see the vocabulary that we used to represent the assets, the vulnerabilities, the reports, and so on.

Copyright: Ontotext New York | Sofia | Basel
Reach out to us: info@ontotext.com

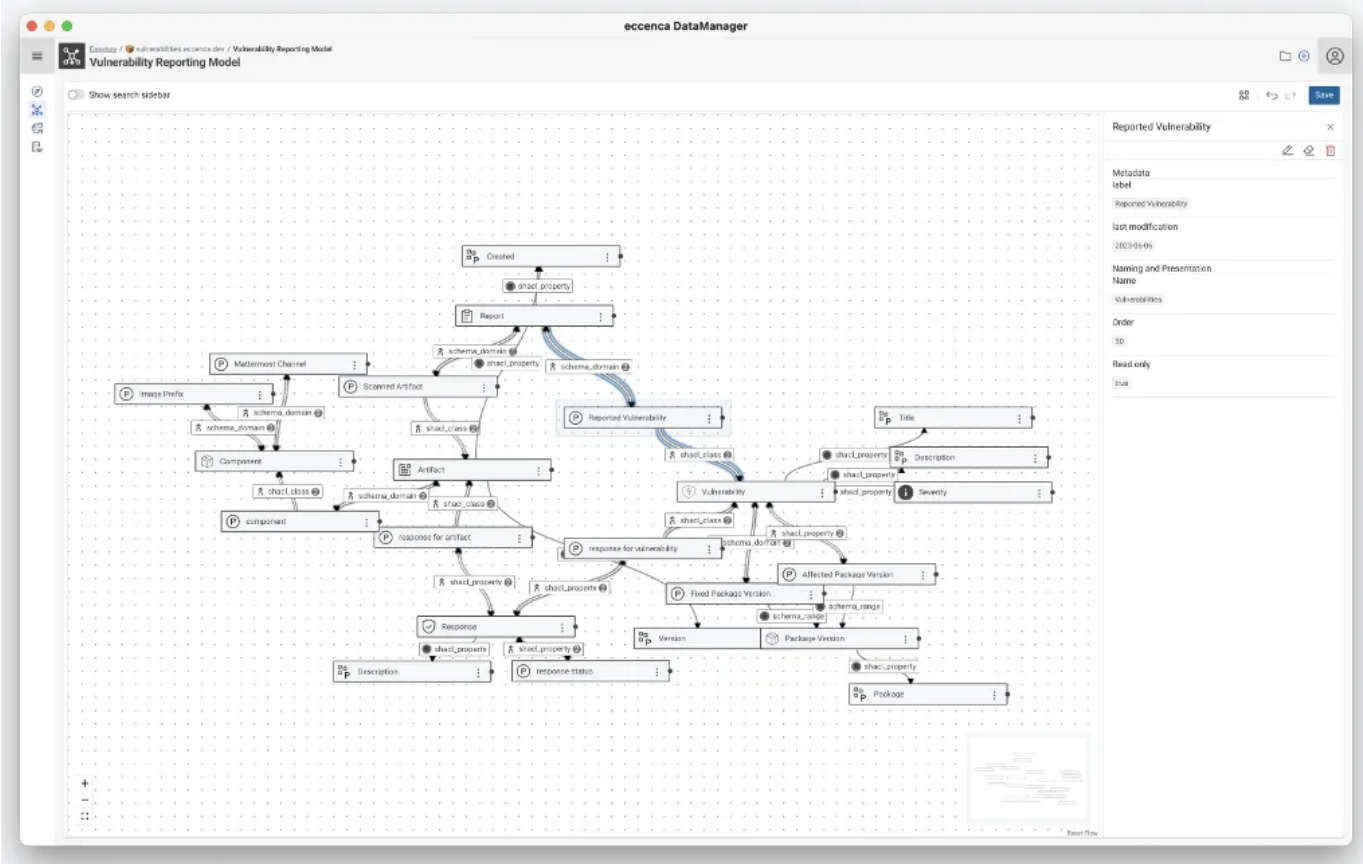

The next figure shows one of the pipelines invoked from the GitLab CI/CD pipeline. This is where the JSON documents from Trivy are processed in order to update the knowledge graph.

Copyright: Ontotext New York | Sofia | Basel
Reach out to us: info@ontotext.com

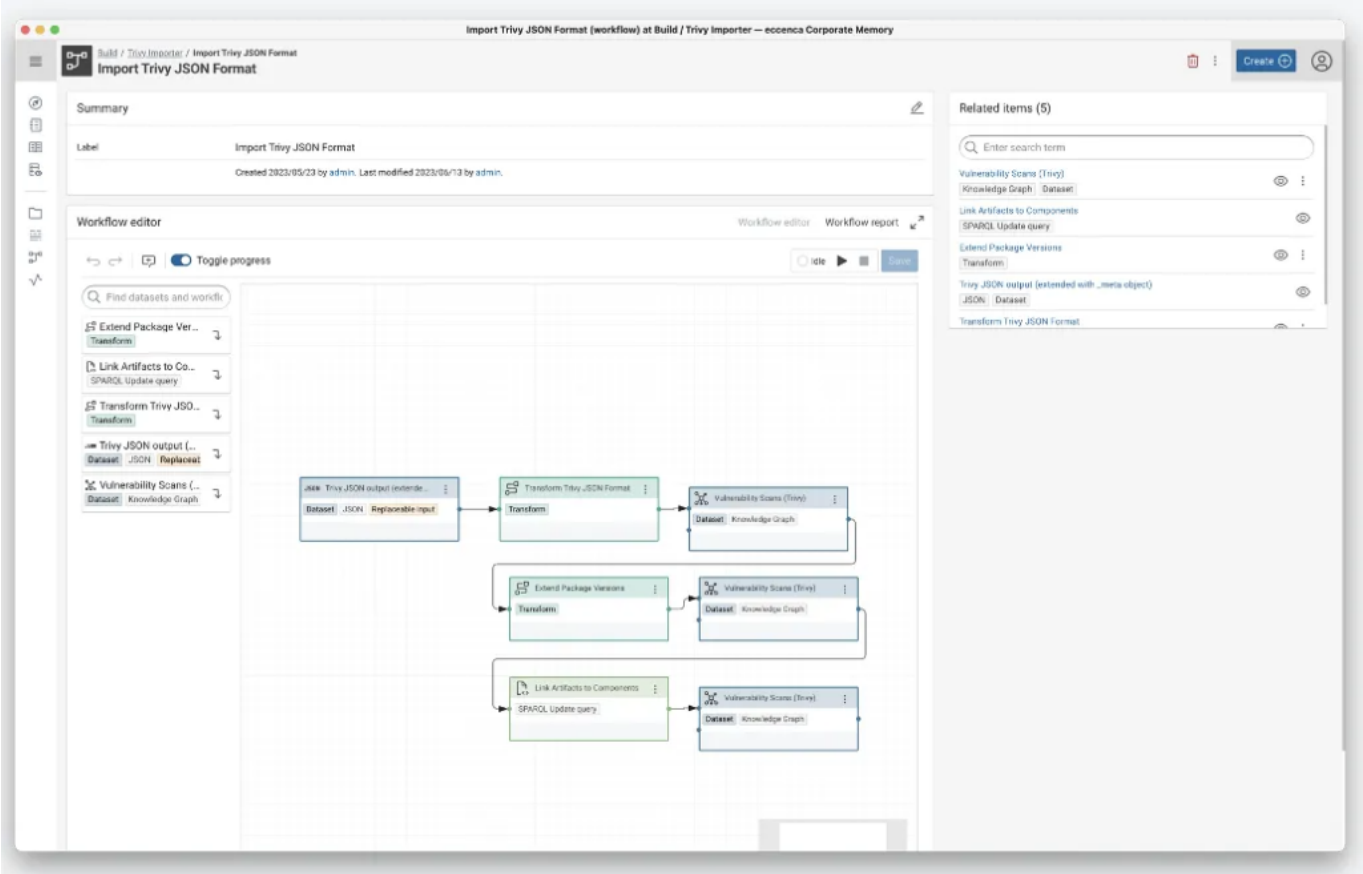

The figure below shows the Corporate Memory based management interface used by experts to enter mitigation strategies and describe the potential risk factors that emerge from the CVEs that we have found.

Copyright: Ontotext New York | Sofia | Basel
Reach out to us: info@ontotext.com

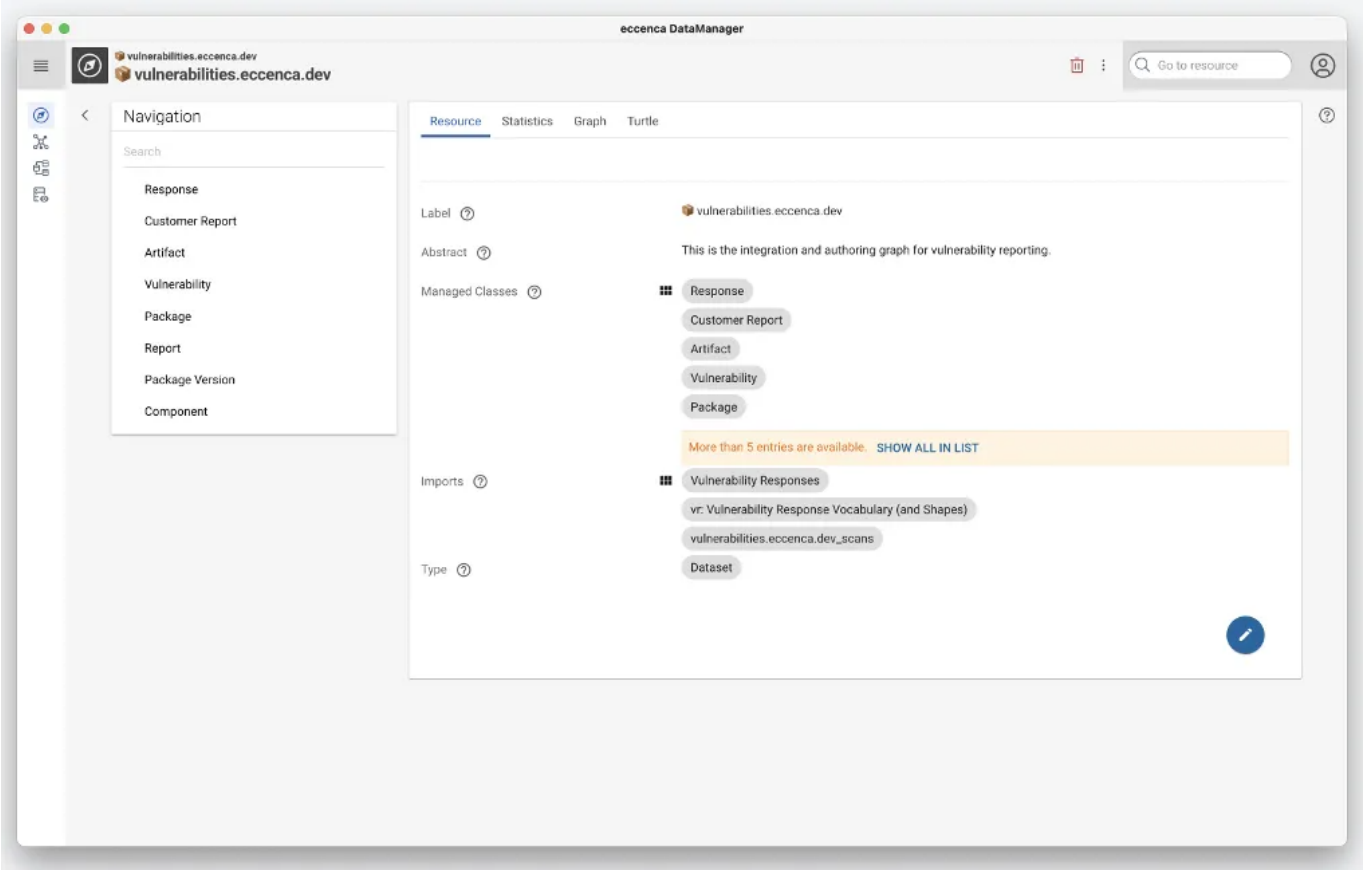

Finally, the messaging process: a second pipeline is invoked, sorting out newly uncommented, unlinked CVEs and signaling them into different Mattermost channels to various teams. These teams are then taken back to our platform, where they can annotate and describe the potential risk that emerges from a particular CVE and what mitigation strategies need to be applied.

Copyright: Ontotext New York | Sofia | Basel
Reach out to us: info@ontotext.com

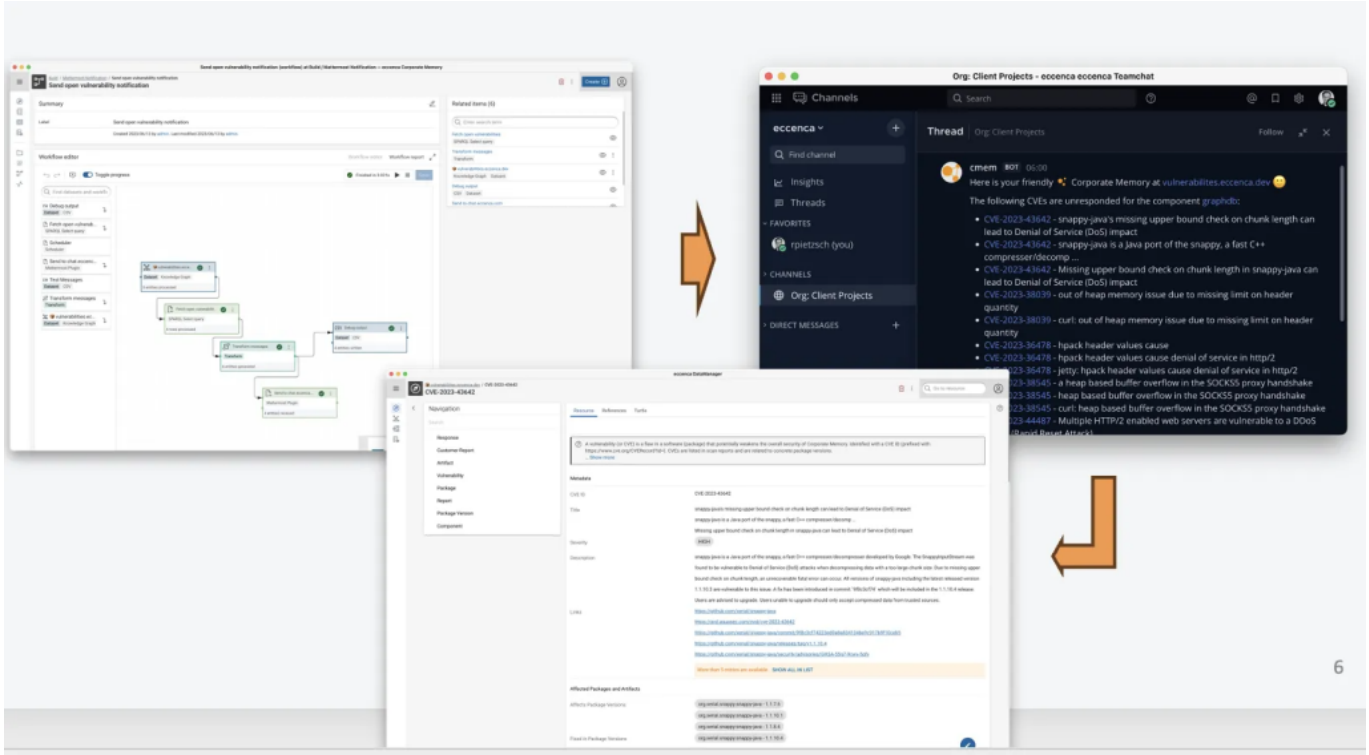

This whole process is used for producing different types of reports and in different granularities as shown in the screenshots below. On the left side, you see a component report, which is the result of a vulnerability report for one specific Docker image. On the right slide is the vulnerability report configuration. The list of all the artifacts scanned in one run of this pipeline.

Copyright: Ontotext New York | Sofia | Basel
Reach out to us: info@ontotext.com

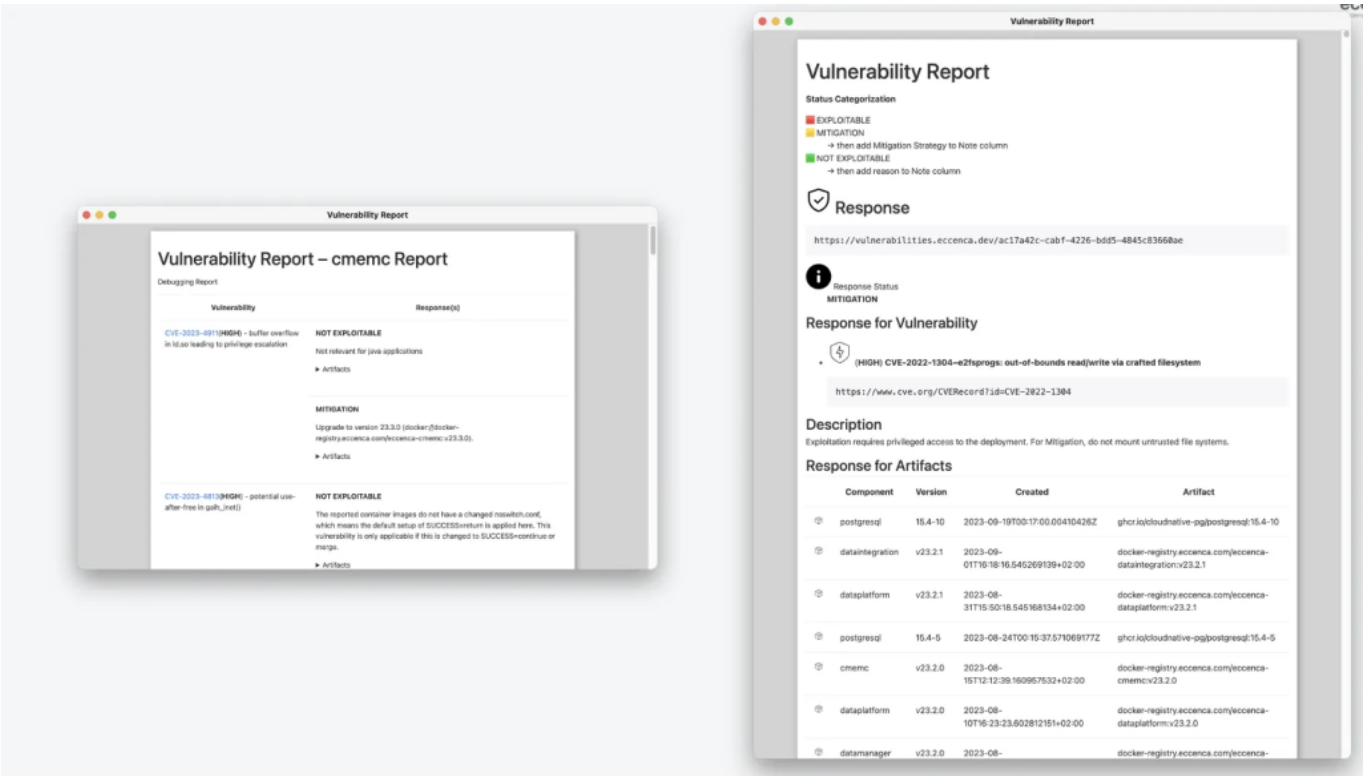

There are also customer-specific reports, which resemble specific customer deployments so that we can give them an early warning and recommend specific actions that need to be applied on their end.

To Wrap It Up

Now, you are not only free to start whatever knowledge graph project on your sandbox, but you can also take this pre-confectioned package as a Cybersecurity solution and apply it to your own needs of CVE and vulnerability management.

Our product, eccenca Corporate Memory, is versatile and useful in many situations. Essentially, it adds value wherever there is a lot of data stored in separate, unconnected systems. Our product not only links this data together but also incorporates crucial expert knowledge. This comprehensive view is essential for making important, forward-thinking business decisions. We partner with Ontotext and use GraphDB as the engine for our knowledge platform. For the business applications of Corporate Memory, we collaborate with

Copyright: Ontotext New York | Sofia | Basel
Reach out to us: info@ontotext.com

partners in the manufacturing, automotive, telecommunications, and life sciences industries.

Do you want to learn more?

Book a Discovery Call

Copyright: Ontotext New York | Sofia | Basel
Reach out to us: info@ontotext.com