Bringing context to operational data with MCP servers

Shane McDowell
Sr. Principal Product Manager, Red Hat Lightspeed

Red Hat

The journey today:

- Our story

  We don’t operate alone.

- Scaling operational data

  One system - easy. Many systems - hard.

- Contextualizing operational data

  Make it make sense.

- Accelerating operational data

  Zoom zoom.





Shane McDowell

- Sr. Principal Product Manager
- RHEL Downloads / RHEL Management / Red Hat Lightspeed (inc. on-prem)
- 14 years with Red Hat

Resident of Michigan (One of us! One of us! One of us!)

Email: smcdowel@redhat.com
LinkedIn: @shanemcdowell1

I am not Dr. Brian Ball.

I am best-equipped to talk about data I know best - Red Hat Lightspeed (fmly Red Hat Insights).

I'll use Claude code today, but that is not a Red Hat endorsement.[1]

1 - Claude code is just one of many LLM clients. It happens to be the one I use due to organizational policy and approval. Please follow your own organization's polcies and InfoSec approvals. Claude, Gemini, Copilot, Goose, ChatGPT, and more all of the time.

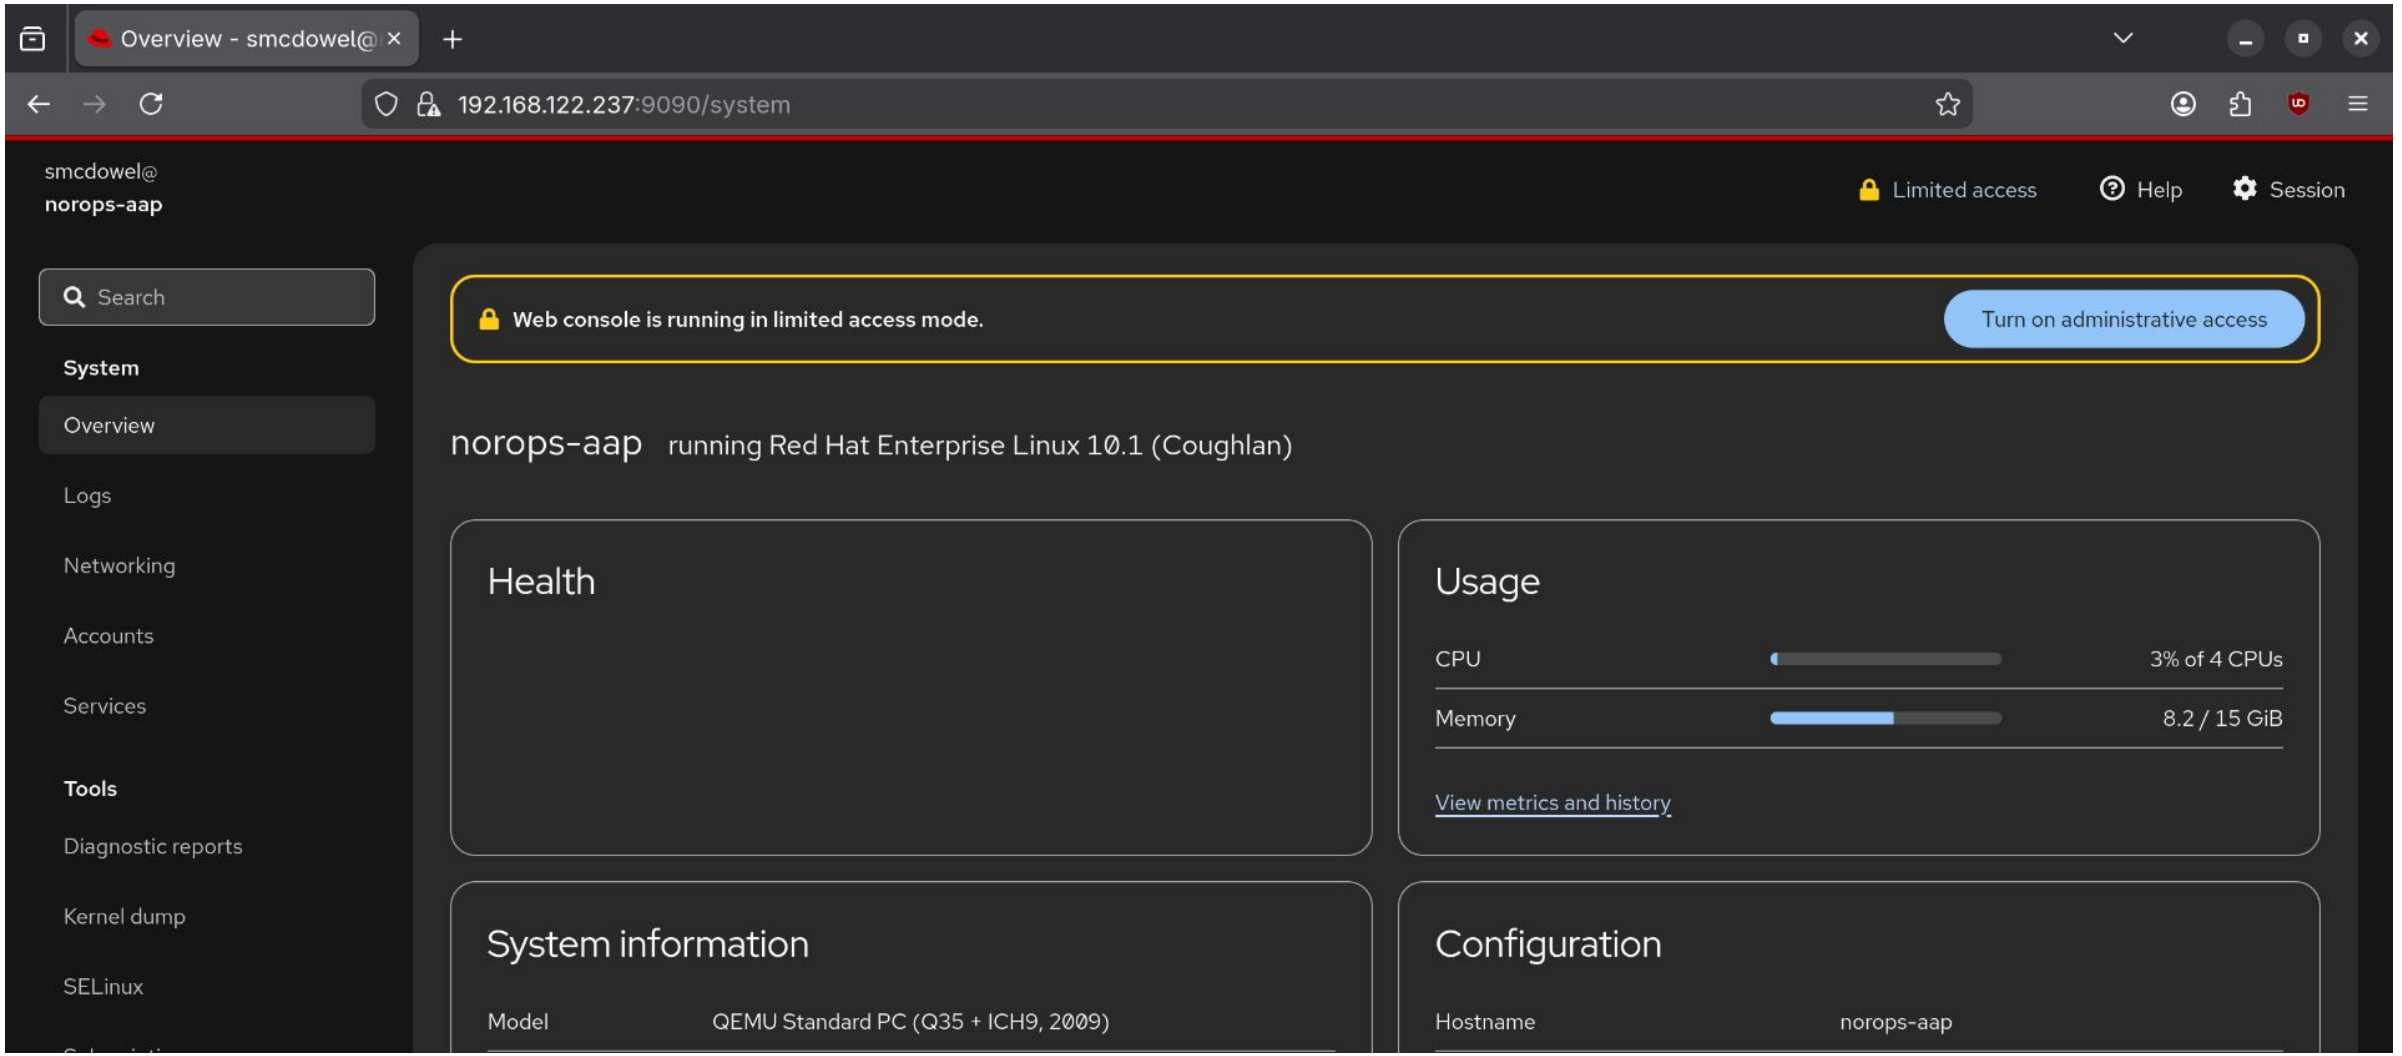


You manage technology, but you don’t manage it alone. You need good operational data for you and your partners to make good operational decisions.

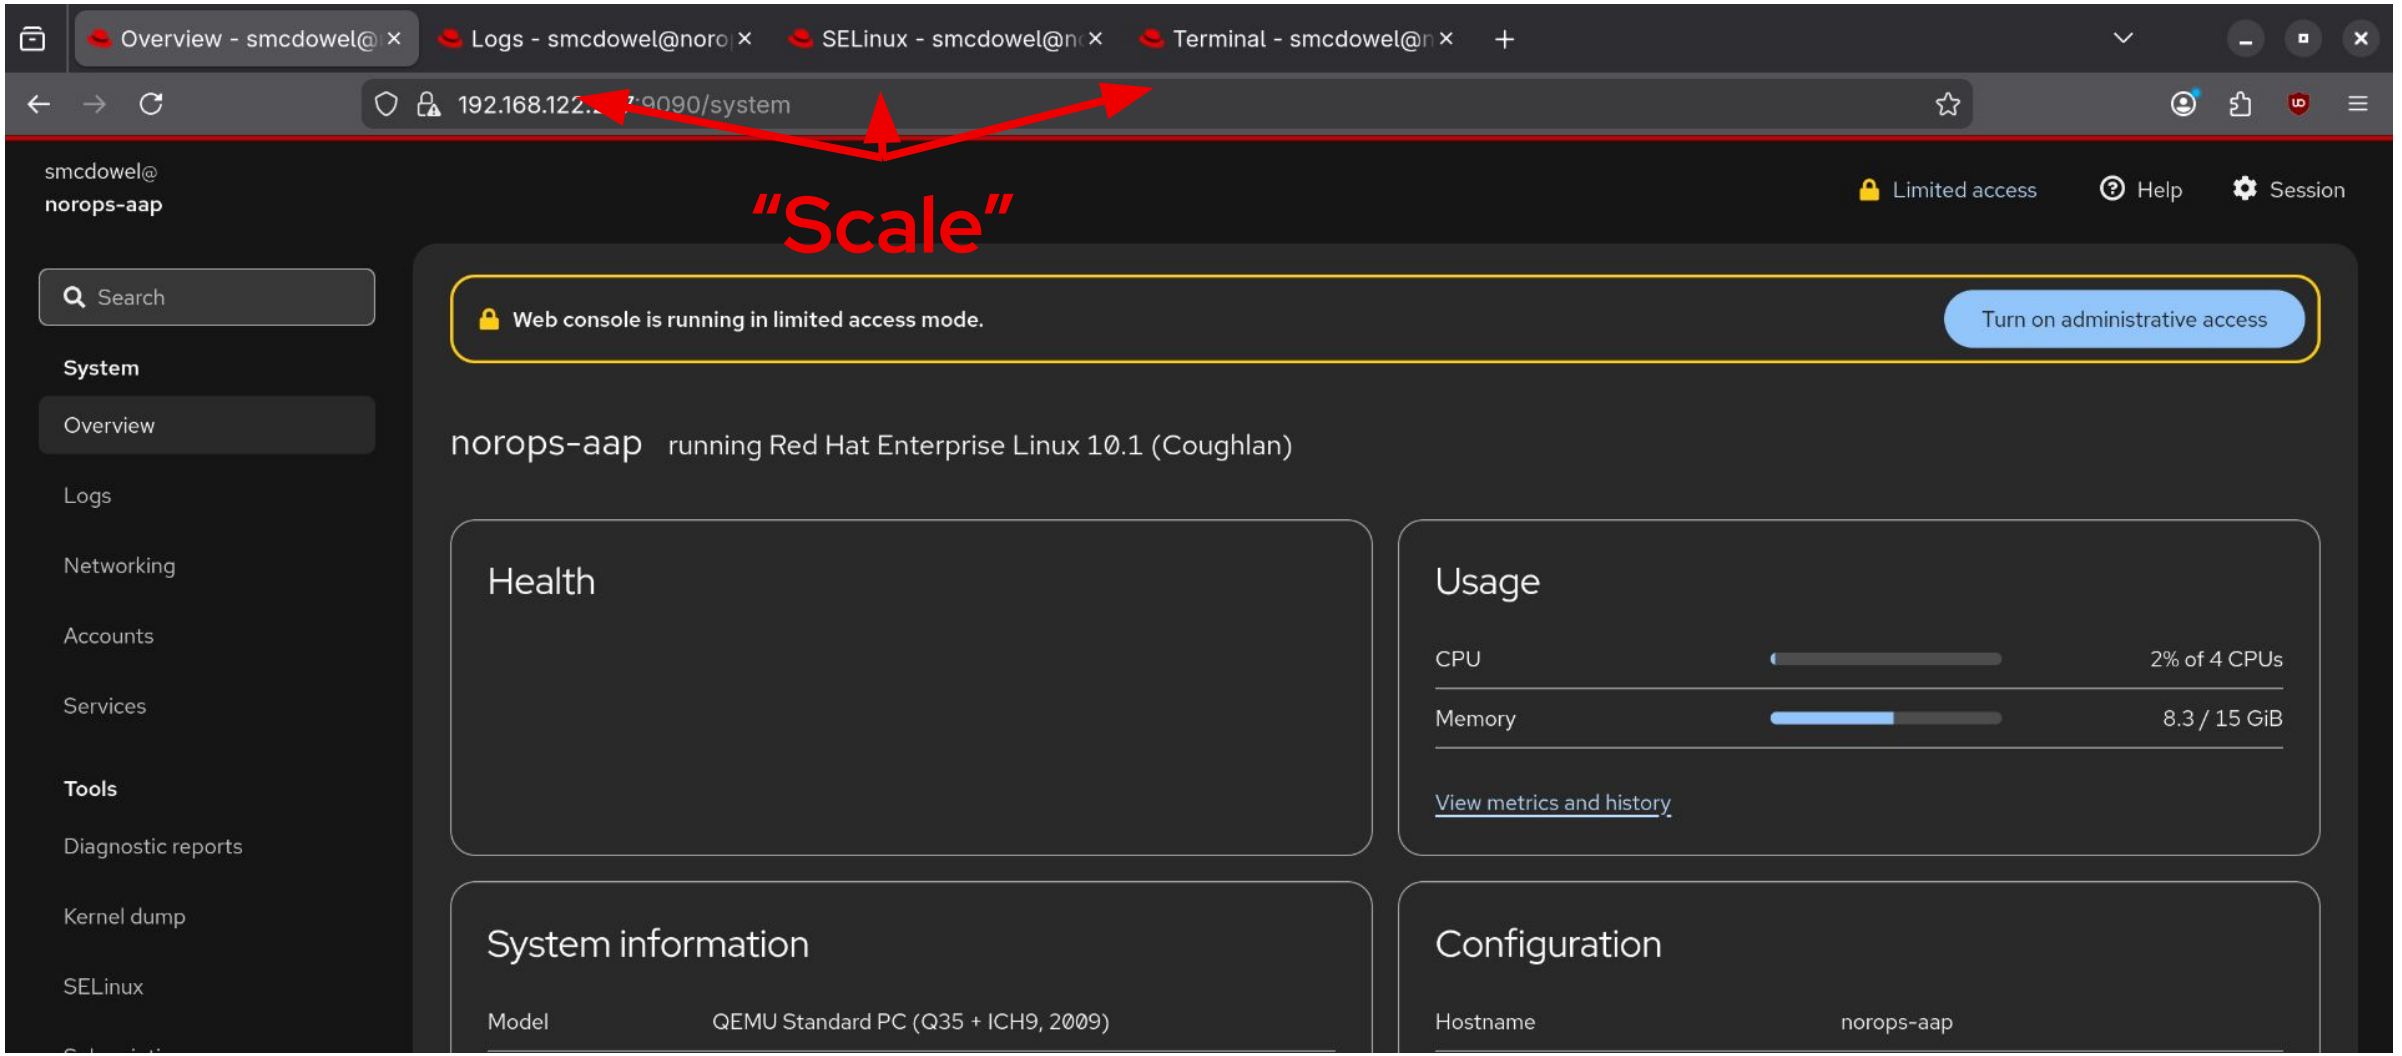
In theory, more data leads to better decisions, but at a large scale, it can be overwhelming without context.



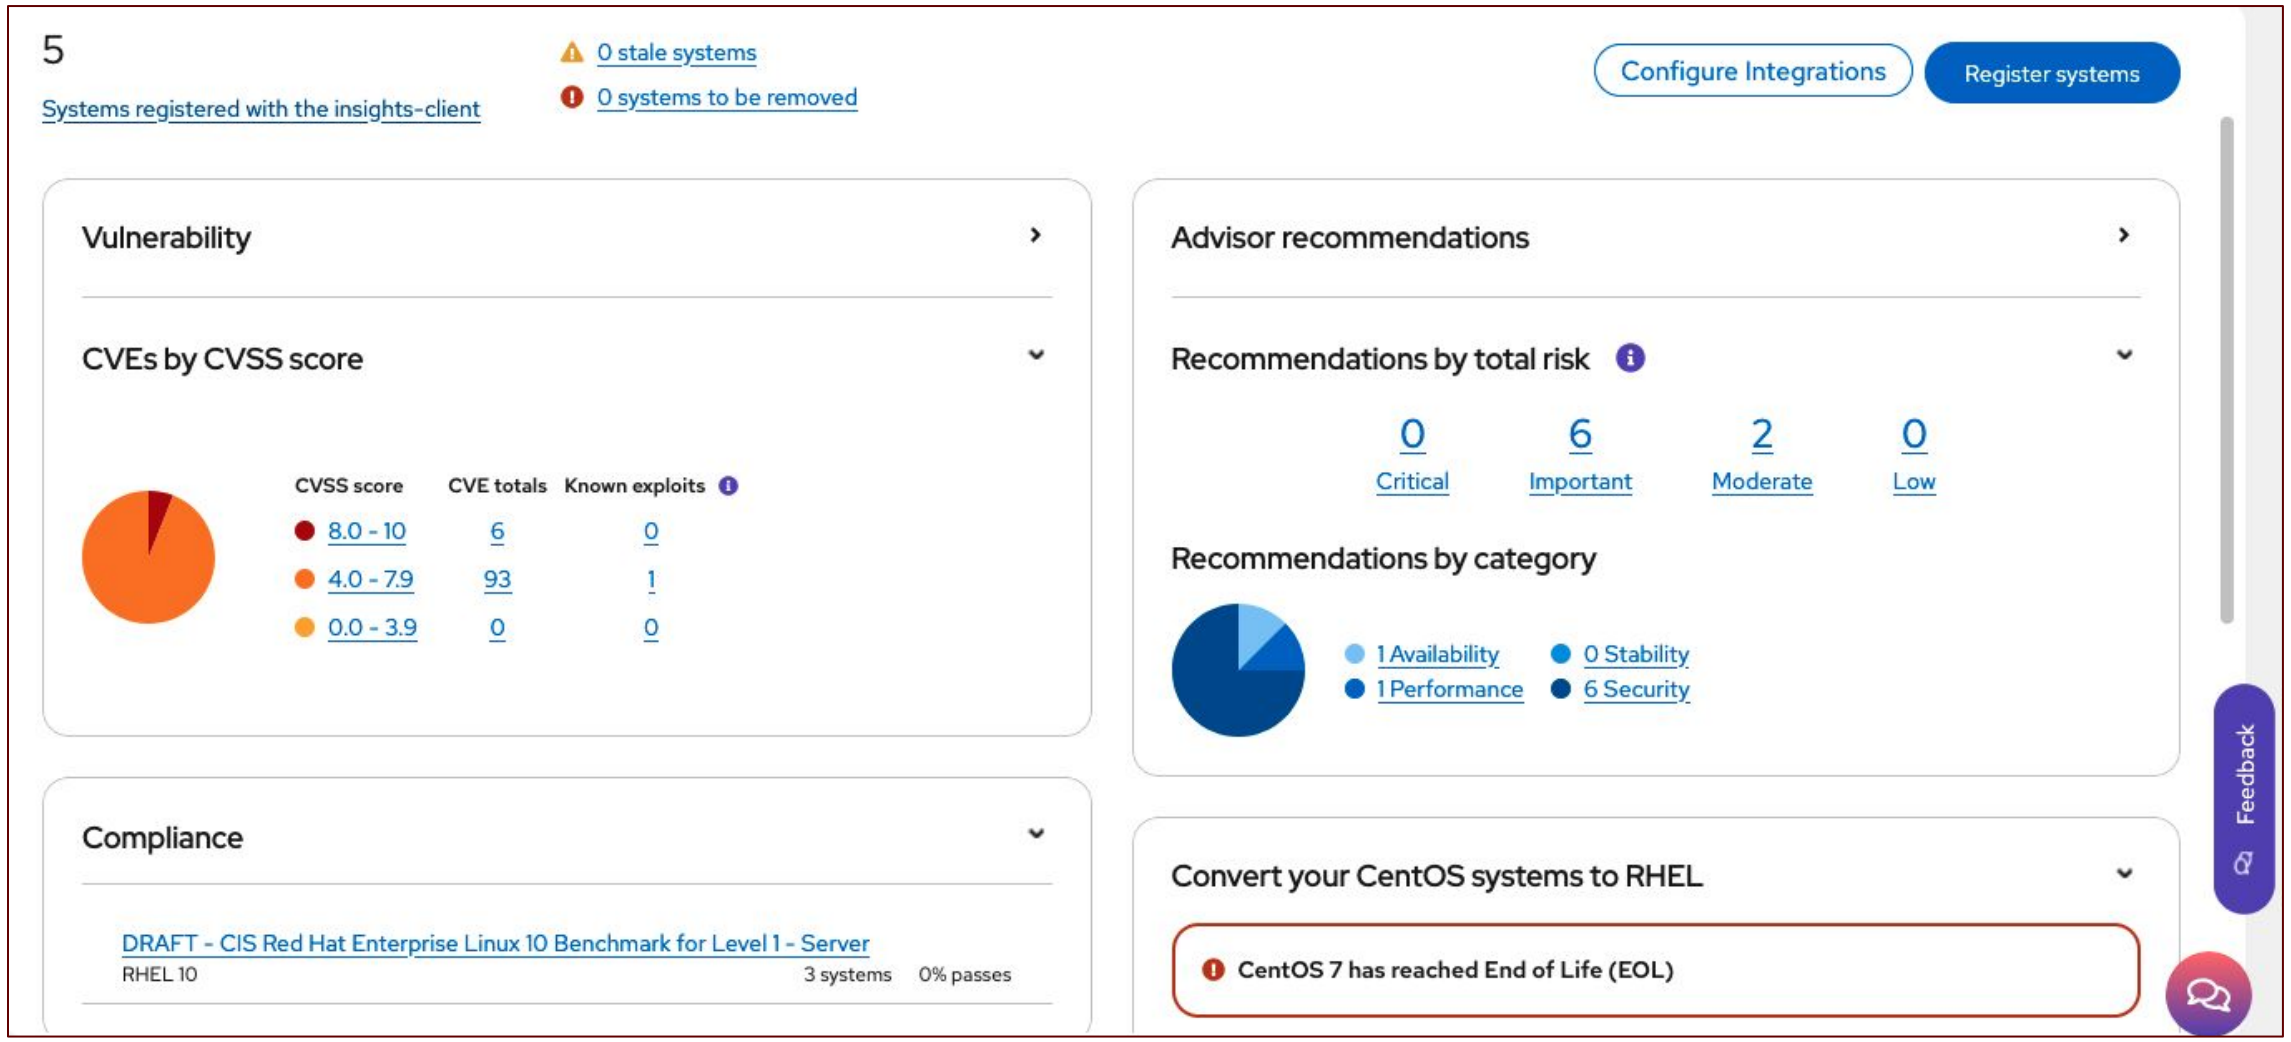
One system, one cluster



Several systems, several clusters



Many systems, many clusters



Many systems via API


```
{
  "hostname": "norops-lab-03",
  "system_id": "d8a3079d-6faf-4bca-9ad5-455e62aa22b5",
  "cves": [
    {
      "cve_id": "CVE-2025-22068",
      "cvss_score": "6.700",
      "advisory_id": "RHSA-2025:22395",
      "advisory_synopsis": "Moderate: kernel security update"
    },
    {
      "cve_id": "CVE-2025-38614",
      "cvss_score": "6.200",
      "advisory_id": "RHSA-2025:21463",
      "advisory_synopsis": "Moderate: kernel security update"
    },
    {
      "cve_id": "CVE-2025-38724",
      "cvss_score": "7.000",
      "advisory_id": "RHSA-2025:22395",
      "advisory_synopsis": "Moderate: kernel security update"
    },
    {
      "cve_id": "CVE-2025-39730",
      "cvss_score": "7.500",
      "advisory_id": "RHSA-2025:21931",
      "advisory_synopsis": "Moderate: kernel security update"
    },
    {
      "cve_id": "CVE-2025-39864",
      "cvss_score": "7.000",
      "advisory_id": "RHSA-2025:21463",
```


We can build some context into our scripts. In this case, I've incorporated our security team's policy that all CVEs must be resolved



What happens when you want to add context from your business operations?



Many systems via API with context


```
    },
    {
      "cve_id": "CVE-2025-39971",
      "published_date": "2025-10-15T00:00:00+00:00",
      "cvss_score": "7.600",
      "advisories": [
        {
          "advisory_id": "RHSA-2025:22395",
          "advisory_synopsis": "Moderate: kernel security update",
          "advisory_published_date": "2025-12-01T08:35:55Z"
        }
      ],
      "affected_systems": [
        "norops-lab-01",
        "norops-lab-02.smcdowel-thinkpadt14sgen2i.rmtusmi.csb",
        "norops-lab-03"
      ],
      "affected_systems_count": 3
    }
  ],
  "compliance_summary": {
    "organizational_standard": "CVEs with errata published more than 30 days ago are out of compliance",
    "total_cves": 12,
    "out_of_compliance_count": 0,
    "compliance_rate": "100.0%",
    "out_of_compliance_cves": []
  },
  "summary": {
    "total_cves_with_advisories": 12,
    "total_systems_affected": 3
  }
}smcdowel@smcdowel-thinkpadt14sgen2i:~/Documents/claude/rh-mir-hug$
```




What if we aren’t sure yet what context is important?

Can we use tools that can infer that context?

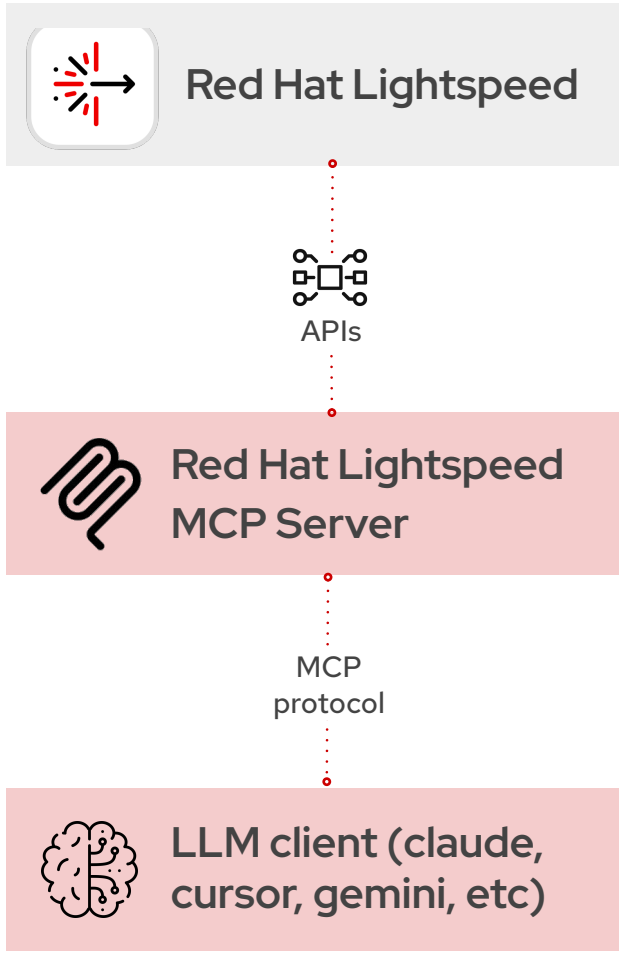


Let’s talk about MCP (Model Context Protocol)

- A standard created by Anthropic for integrating AI applications with various data sources
- This is what drives talk of “agentic AI” - agents (MCP servers) integrating and communicating context
- At its very most basic level, it can infer context from API endpoints



Anthropic MCP - https://www.anthropic.com/news/model-context-protocol



- Available now as Technical Preview (source on github)
- Deployed as a standalone container (e.g. podman).
- Example of integrations with VSCode, Cursor, Gemini CLI, Claude Desktop
- Using token-based authentication via Service Account (honoring RBAC permissions)

> summarize the recommendations for my RHEL environment

```
● Summary of Advisor Recommendations for Your RHEL Environment

  Overview Statistics

  - Total recommendations: 7 active issues
  - Risk distribution:
    - High risk (3): 5 recommendations
    - Medium risk (2): 2 recommendations
  - Categories:
    - Security: 6 recommendations
    - Performance: 1 recommendation

  Top Recommendations (by total risk)

  High Risk Security Issues (5)

  1. Incorrect permissions on sensitive files (3 systems affected)
    - Rule: hardening_passwd_perms|HARDENING_PASSWD_PERMS
    - Impact: Authentication Bypass
    - Issue: Password/shadow files have incorrect permissions/ownership
    - https://console.redhat.com/insights/advisor/recommendations/hardening_passwd_perms%7CHARDENING_PASSWD_PERMS
  2. OpenSSH config permissions (3 systems affected)
    - Rule: hardening_ssh_config_perms|OPENSSH_HARDENING_CONFIG_PERMS
    - Impact: Decreased Security
    - Issue: SSH configuration not following security best practices
    - https://console.redhat.com/insights/advisor/recommendations/hardening_ssh_config_perms%7COPENSSH_HARDENING_CONFIG_PERMS
  3. OpenSSH PermitEmptyPasswords (2 systems affected)
    - Rule: hardening_ssh_empty_passwords|OPENSSH_HARDENING_EMPTY_PASSWORDS
    - Impact: Decreased Security
    - Issue: Empty passwords permitted in SSH configuration
  4. OpenSSH LoginGraceTime (3 systems affected)
    - Rule: hardening_ssh_grace_time|OPENSSH_HARDENING_GRACE_TIME
```



> briefly summarize the health of my environment considering the inventory, recommendations, and security vulnerabilities that have errata available

```
● NorOps Lab Environment Health Summary

  Inventory

  ✅ 4 RHEL systems - All active and reporting normally
  - norops-aap (16GB RAM)
  - norops-lab-01, norops-lab-02, norops-lab-03 (4GB RAM each)

  Security Vulnerabilities

  ✅ 0 Critical/Important CVEs with errata available - Excellent!

  Advisor Recommendations

  ⚠️7 active issues requiring attention:

  Security (6 issues)
  - 5 High Risk: SSH hardening concerns across 3 systems
    - Insecure SSH configurations (empty passwords permitted, weak timeouts, high auth retries)
    - File permission issues on authentication files (passwd/shadow)
  - 1 Medium Risk: Auditd service not enabled on 2 systems

  Performance (1 issue)
  - Medium Risk: PostgreSQL running without optimal tuned profile on 3 systems

  Overall Assessment

  Health: GOOD - No critical vulnerabilities with patches available, but security hardening recommended. Primary concerns are SSH
  configuration weaknesses and missing audit logging on some systems. All have automated Ansible remediation available.

  Immediate Priority: Address the 5 high-risk SSH and file permission issues to reduce attack surface.
```

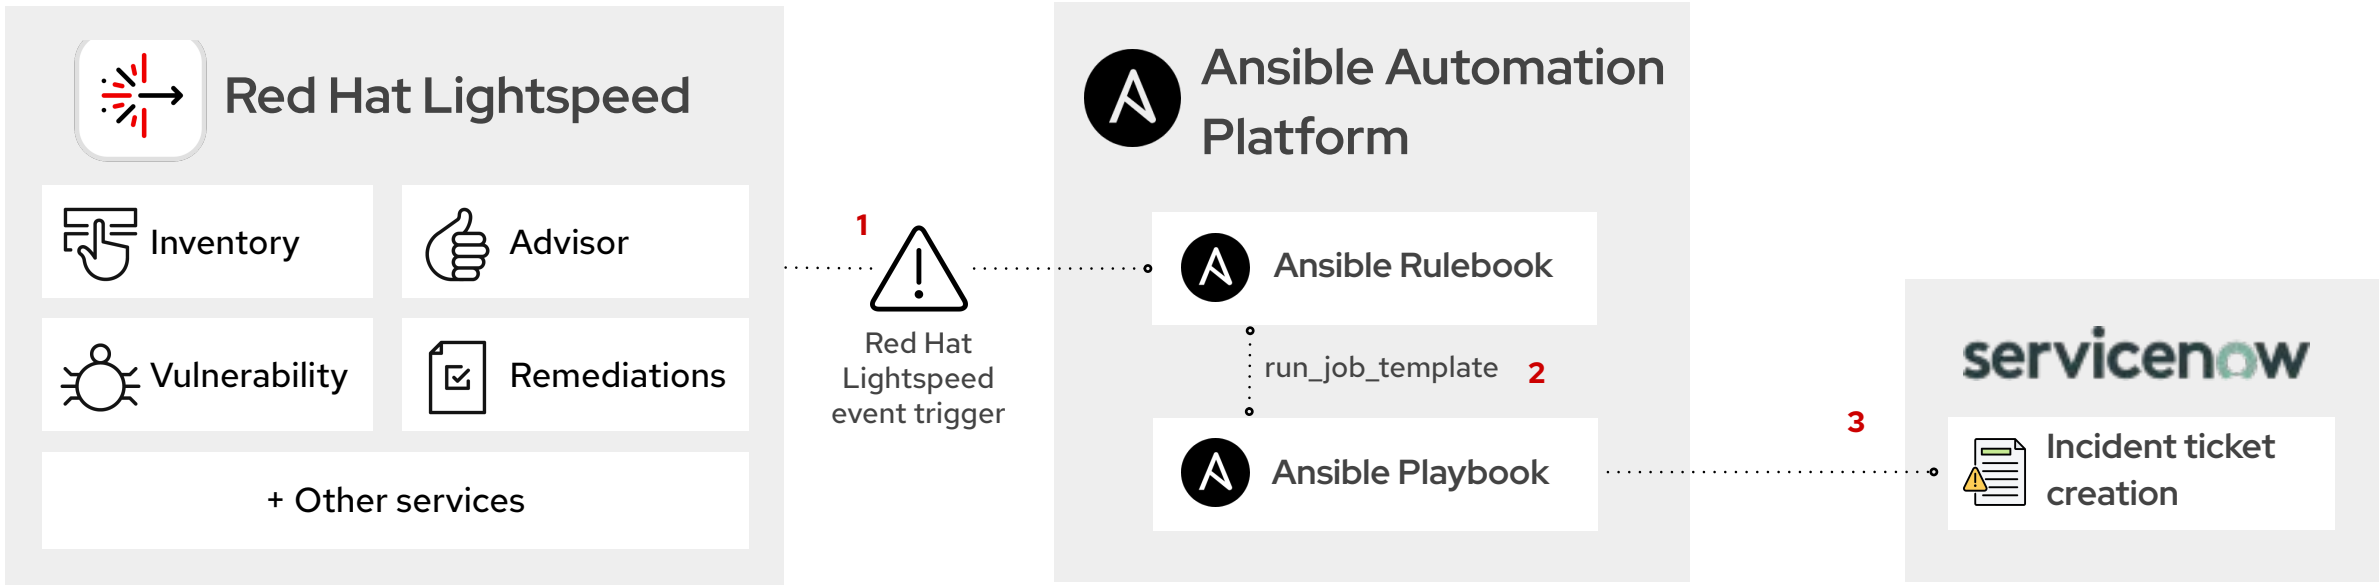


Instead of an individual with a command prompt, what if the bots could talk to each other?

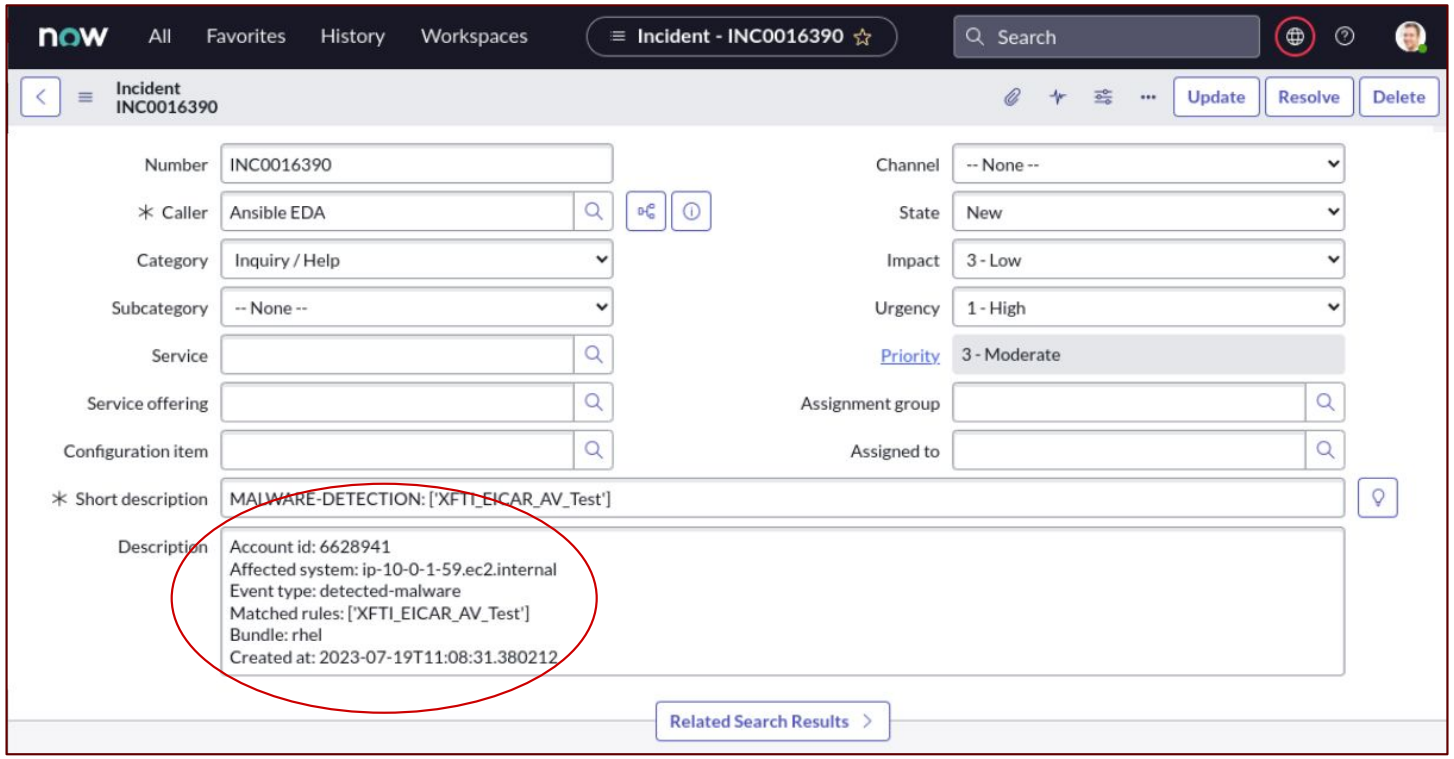
(No, it doesn't result in Skynet)



Example current EDA workflow for Red Hat Lightspeed



Resulting ticket from EDA workflow



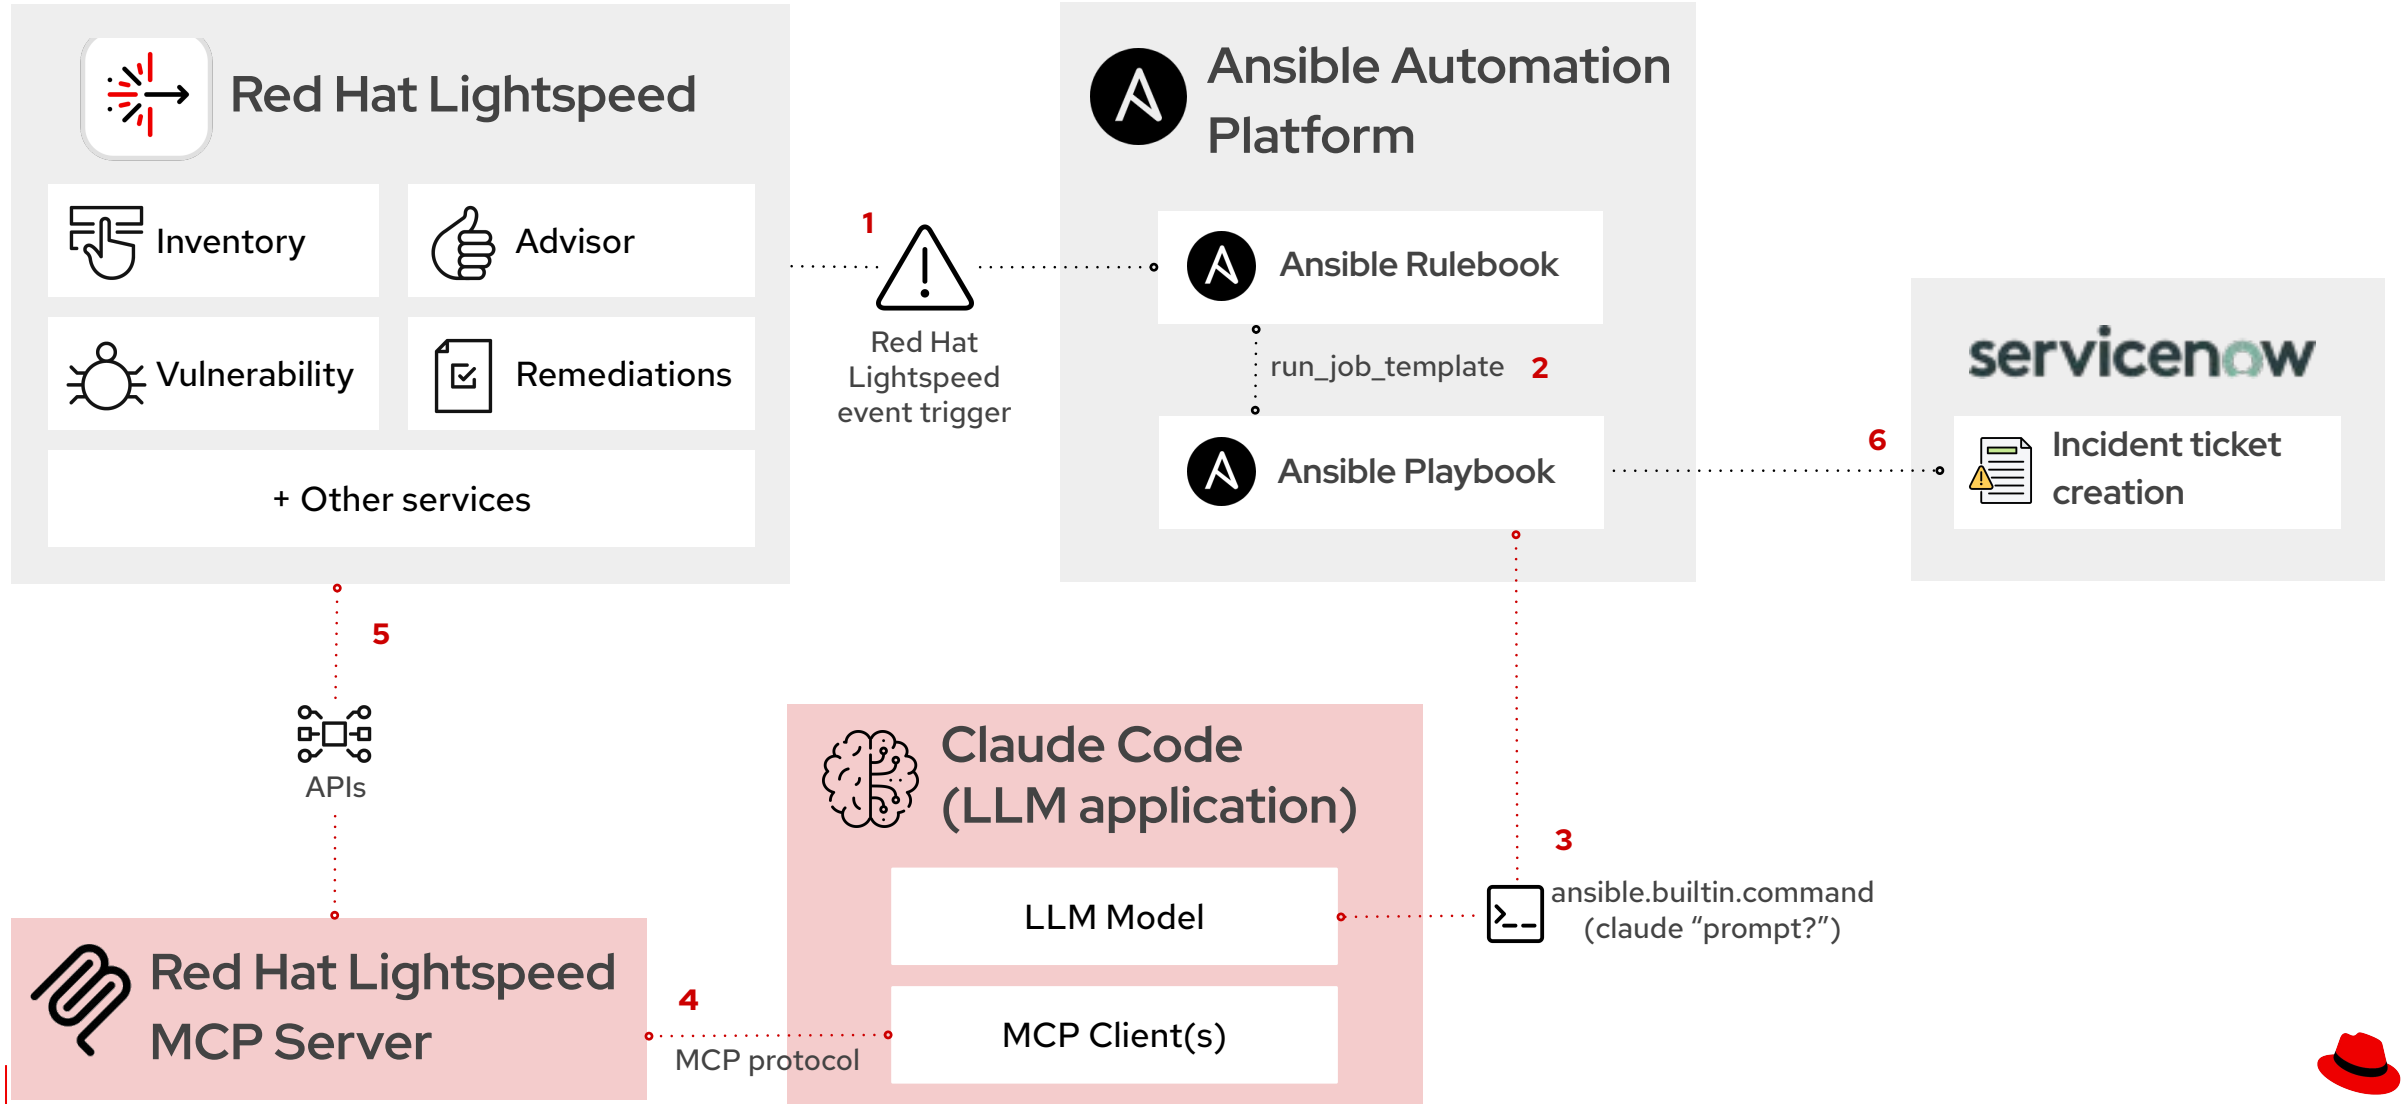
- Missing context
- No severity or priority
- No probable cause
- No recommended action
- Poor readability

What if we could help Operations by providing additional and relevant knowledge in the ticket?



Let’s have Ansible ask an LLM.

LLM-enhanced EDA workflow with Red Hat Lightspeed



Example Ansible task snippet

```
- name: Enrich incident using Claude
  command: >
    claude --output-format json
    "Given this event data, summarize and enhance the ServiceNow
ticket with context, impact, and recommended actions:
    {{ lightspeed_event }}"
  register: llm_out

- set_fact:
    enriched: "{{ llm_out.stdout | from_json }}"

- name: Create ServiceNow ticket
  servicenow.itsm.incident:
    short_description: "{{ enriched.summary }}"
    description: "{{ enriched.details }}"
```



Original content

```
Short description: MALWARE-DETECTION:
['XFTI_EICAR_AV_Test']

Description:
Account id: 6628941
Affected system:
ip-10-0-1-59.ec2.internal
Event type: detected-malware
Matched rules: ['XFTI_EICAR_AV_Test']
Bundle: rhel
Created at: 2023-07-19T11:08:31.380212
```

LLM enhanced

```
**Short description:**
[Security Alert] EICAR Test File Detected on RHEL Host
(ip-10-0-1-59.ec2.internal)

**Description:**
A malware detection event has been reported for account 6628941
on host ip-10-0-1-59.ec2.internal.
The rule matched the EICAR antivirus test signature
(XFTI_EICAR_AV_Test), which is a harmless sample used to
validate antivirus functionality.

**Summary:**
- Detection type: Test signature (EICAR)
- Severity: Informational (no real threat)
- Bundle: rhel
- Timestamp: 2023-07-19T11:08:31 UTC

**Impact:**
No production impact expected. Detection indicates that the
host's malware detection engine is functioning correctly.

**Recommended Action:**
1. Confirm that the alert corresponds to an intentional EICAR
test.
2. If unintentional, validate AV rules and review recent file
transfers.
3. Document this as a successful antivirus validation event.

**LLM Insight:**
This event represents a benign antivirus test file detection
rather than a real infection. The system appears to be operating
correctly with up-to-date detection rules.
```

```
**Short description:**
[Security Alert] EICAR Test File Detected on RHEL Host
(ip-10-0-1-59.ec2.internal)

**Description:**
A malware detection event has been reported for account 6628941
on host ip-10-0-1-59.ec2.internal.
The rule matched the EICAR antivirus test signature
(XFTI_EICAR_AV_Test), which is a harmless sample used to
validate antivirus functionality.

**Summary:**
- Detection type: Test signature (EICAR)
- Severity: Informational (no real threat)
- Bundle: rhel
- Timestamp: 2023-07-19T11:08:31 UTC

**Impact:**
No production impact expected. Detection indicates that the
host's malware detection engine is functioning correctly.

**Recommended Action:**
1. Confirm that the alert corresponds to an intentional EICAR
test.
2. If unintentional, validate AV rules and review recent file
transfers.
3. Document this as a successful antivirus validation event.

**LLM Insight:**
This event represents a benign antivirus test file detection
rather than a real infection. The system appears to be operating
correctly with up-to-date detection rules.
```

Ticket enrichment highlights

- Human-readable context
  Explained EICAR as a test signature
- Severity classification
  Labeled as “Informational”
- Impact assessment
  Clarified no threat / test event
- Recommended actions
  Added 3 specific next steps
- Summarized technical fields
  Turned data into narrative



Operational data leads to good decisions, but we need tools to help us with context, especially when getting support from others.

MCP servers, LLMs, and agentic workflows can provide the context and inference needed to accelerate operations.



Thank you

Red Hat is the world's leading provider of enterprise open source software solutions. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500.


linkedin.com/company/red-hat


youtube.com/user/RedHatVideos


facebook.com/redhatinc


x.com/RedHat

Red Hat