Security and Distributed Cloud Services Release Notes – May 2023

Welcome to the product release notes from F5® Distributed Cloud Services. The product team will provide additional details on key features and enhancements in every release each month.

Web Application and API Protection

API Endpoints Risk Score

We are announcing the release of a new Risk Score feature, which provides customers with a comprehensive measure of the risk associated with their API endpoints. The Risk Score is calculated using a variety of techniques, such as vulnerability discovery, attack impact, business value, attack likelihood, and mitigating controls, to help customers evaluate the potential impact of vulnerabilities or threats to an API endpoint and prioritize efforts to mitigate those risks. Customers can view the content of the Risk Score in the Security Posture tab that appears in the Endpoint Details of each API endpoint, with instructions and evidence for each vulnerability.

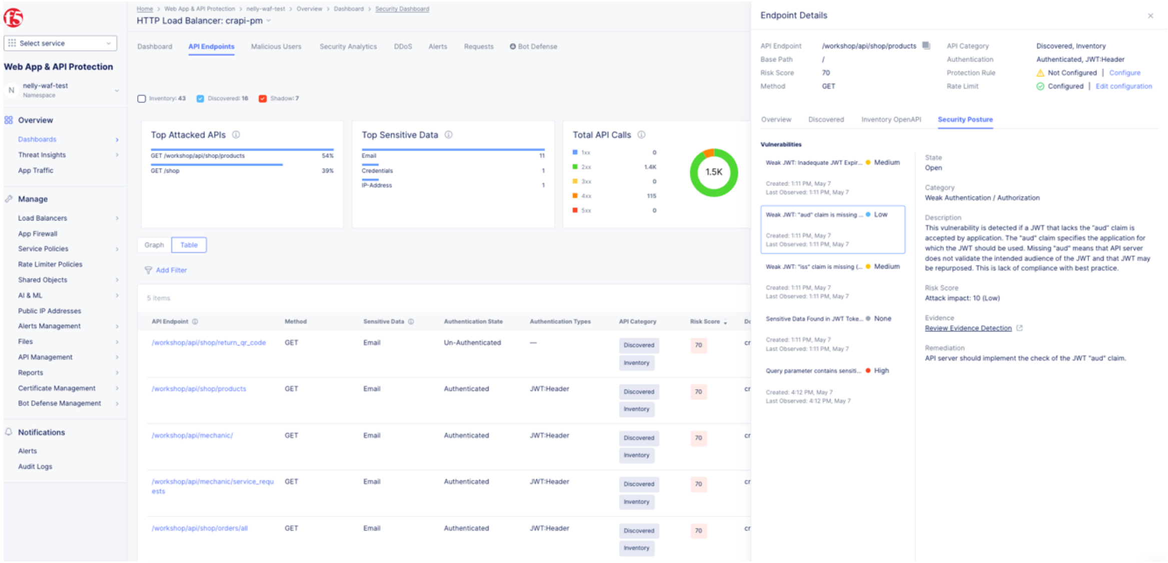

☐ Inventory: 43 ☑ Discovered: 16 ☑ Shadow: 7

Top Attacked APIs

GET /workshop/api/shop/products	54%
GET /shop	39%

Top Sensitive Data

Email	11
Credentials	1
IP-Address	1

Total API Calls

🟦 1xx	0
🟩 2xx	1.4K
🟨 3xx	0
🟧 4xx	115
🟥 5xx	0

1.5K

5 items

API Endpoint	Method	Sensitive Data	Authentication State	Authentication Types	API Category	Risk Score
/workshop/api/shop/return_qr_code	GET	Email	Un-Authenticated	—	Discovered Inventory	70
/workshop/api/shop/products	GET	Email	Authenticated	JWT:Header	Discovered Inventory	70
/workshop/api/mechanic/	GET	Email	Authenticated	JWT:Header	Discovered Inventory	70
/workshop/api/mechanic/service_requests	GET	Email	Authenticated	JWT:Header	Discovered Inventory	70
/workshop/api/shop/orders/all	GET	Email	Authenticated	JWT:Header	Discovered Inventory	70

Endpoint Details

API Endpoint	/workshop/api/shop/products	API Category	Discovered, Inventory
Base Path	/	Authentication	Authenticated, JWT:Header
Risk Score	70	Protection Rule	⚠ Not Configured | Configure
Method	GET	Rate Limit	✓ Configured | Edit configuration

Vulnerabilities

- Weak JWT: Inadequate JWT Expir... ● Medium — Created: 1:11 PM, May 7 Last Observed: 1:11 PM, May 7
- Weak JWT: "aud" claim is missing ... ● Low — Created: 1:11 PM, May 7 Last Observed: 1:11 PM, May 7
- Weak JWT: "iss" claim is missing (... ● Medium — Created: 1:11 PM, May 7 Last Observed: 1:11 PM, May 7
- Sensitive Data Found in JWT Toke... ● None — Created: 1:11 PM, May 7 Last Observed: 1:11 PM, May 7
- Query parameter contains sensiti... ● High — Created: 4:12 PM, May 7 Last Observed: 4:12 PM, May 7

State
Open

Category
Weak Authentication / Authorization

Description
This vulnerability is detected if a JWT that lacks the "aud" claim is accepted by application. The "aud" claim specifies the application for which the JWT should be used. Missing "aud" means that API server does not validate the intended audience of the JWT and that JWT may be repurposed. This is lack of compliance with best practice.

Risk Score
Attack impact: 10 (Low)

Evidence
Review Evidence Detection

Remediation
API server should implement the check of the JWT "aud" claim.

Figure 1: Risk scoring for API endpoints delivers insights into details of vulnerable API endpoints, streamlining remediation efforts.

API Spec Enforcement

Many of the threats and issues within the OWASP API Security Top 10 are triggered by the lack of input or output validation of APIs. An API-native, positive security approach is required to protect APIs from such issues. With this new API spec enforcement functionality users can easily create an allowlist of the characteristics for valid, allowed requests. These characteristics are used to validate input and output data for things like data type, min or max length, permitted characters, or valid value ranges. This allows users to more easily check if traffic complies with an API schema and block if not. Users can configure validation for individual API endpoints, API groups, or base paths within the API definition.

Mutual TLS Authentication: Send Client Certificate Details to Origin Servers

Mutual TLS now supports the ability to send client certificate details to origin servers in x-forwarded-client-cert (XFCC) request headers. Some origin servers may need client certificate information for audit purposes or may want to issue tokens or cookies to a client certificate. Users can choose to send the entire client certificate or specific fields via the XFCC header.

New Security Incidents Functionality

The Security Incidents functionality simplifies the investigation of attacks by grouping thousands of events into a few incidents based on context and common characteristics. The incidents for the HTTP load balancer are seen in the Incidents tab of the Security Analytics page. This is aimed at making the investigation of application security events easier so IT and operations teams can respond quickly and decisively.

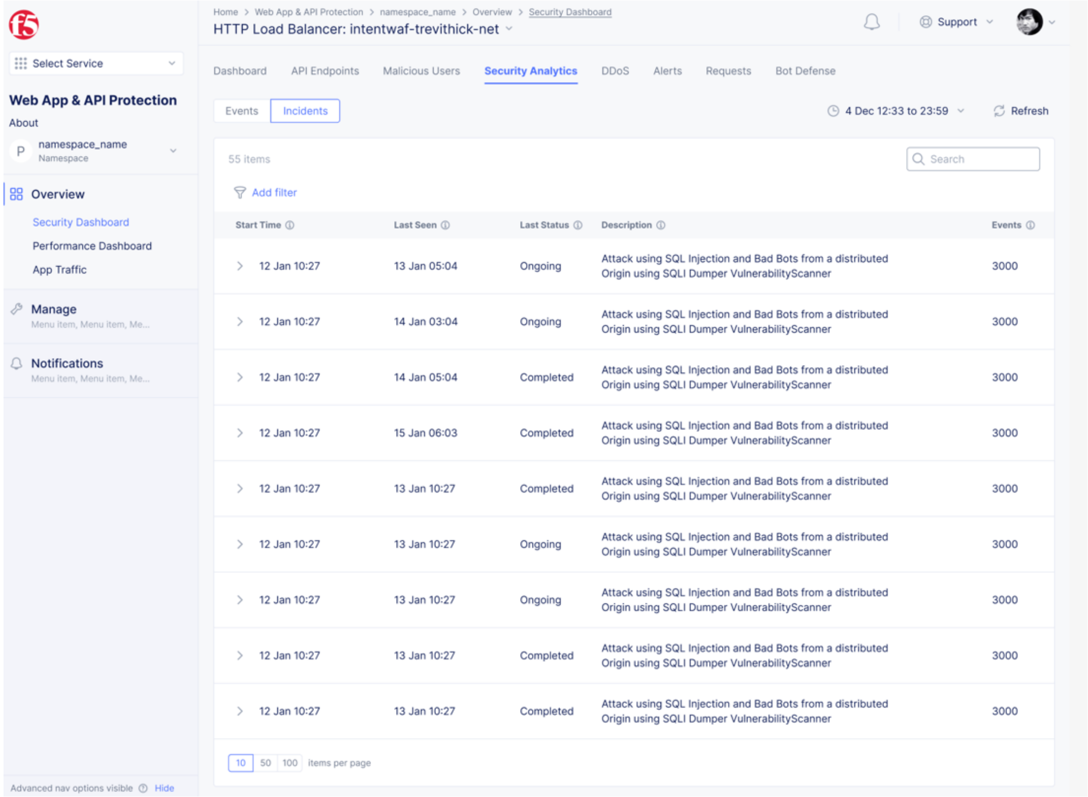

Start Time	Last Seen	Last Status	Description	Events
12 Jan 10:27	13 Jan 05:04	Ongoing	Attack using SQL Injection and Bad Bots from a distributed Origin using SQLI Dumper VulnerabilityScanner	3000
12 Jan 10:27	14 Jan 03:04	Ongoing	Attack using SQL Injection and Bad Bots from a distributed Origin using SQLI Dumper VulnerabilityScanner	3000
12 Jan 10:27	14 Jan 05:04	Completed	Attack using SQL Injection and Bad Bots from a distributed Origin using SQLI Dumper VulnerabilityScanner	3000
12 Jan 10:27	15 Jan 06:03	Completed	Attack using SQL Injection and Bad Bots from a distributed Origin using SQLI Dumper VulnerabilityScanner	3000
12 Jan 10:27	13 Jan 10:27	Completed	Attack using SQL Injection and Bad Bots from a distributed Origin using SQLI Dumper VulnerabilityScanner	3000
12 Jan 10:27	13 Jan 10:27	Ongoing	Attack using SQL Injection and Bad Bots from a distributed Origin using SQLI Dumper VulnerabilityScanner	3000
12 Jan 10:27	13 Jan 10:27	Ongoing	Attack using SQL Injection and Bad Bots from a distributed Origin using SQLI Dumper VulnerabilityScanner	3000
12 Jan 10:27	13 Jan 10:27	Completed	Attack using SQL Injection and Bad Bots from a distributed Origin using SQLI Dumper VulnerabilityScanner	3000
12 Jan 10:27	13 Jan 10:27	Completed	Attack using SQL Injection and Bad Bots from a distributed Origin using SQLI Dumper VulnerabilityScanner	3000

Figure 2: New Incidents view, available within the Security Analytics tab for each HTTP load balancer.

Streamlining of Bot Defense Configuration for WAAP Customers

This release simplifies the configuration and management of Bot Defense for WAAP customers. Users can now edit Bot Defense configurations within the F5 Distributed Cloud Bot Defense service tile or within the HTTP LB configuration pane, enabling customers to manage their configuration from whatever view better fits their way of working. For more information, see Bot Defense end-user documentation for new option(s) on configuring and managing the Bot Defense service.

F5 Distributed Cloud Console, DDoS and Transit Services Self-Service Prefix Advertising

Customers that currently operate and have visibility into the DDoS and Transit Services on the F5 Distributed Cloud Platform, will now be able to control the route announcement or suppression for the internet prefixes under their management, via the service Console or API. The functionality to control the route announcement complements the ability for customers to add, modify, or delete Autonomous System Numbers (ASN), Internet netblocks (IPv4/IPv6), and GRE tunnel configurations. Control for the announcement of an IP netblock through F5 Distributed Cloud carriers allows Managed DDoS Always Available customers to quickly and easily change routes to mitigate an attack. Once the network traffic is routed through the F5 Distributed Cloud DDoS mitigation service, customers will obtain immediate protection against volumetric DDoS attacks. Customers may also continue to collaborate with the F5 Security Operations Center as required to control the route announcement or suppression.

Application and Network Performance

Multi-Cloud Networking

New Multi-Cloud Network Connect Dashboard

The Multi-Cloud Network Connect service now has refreshed dashboards. This includes rich new insights for network operators who can observe and act on their multi-cloud network components with dashboard insights focused on networking, performance, network security, and site management.

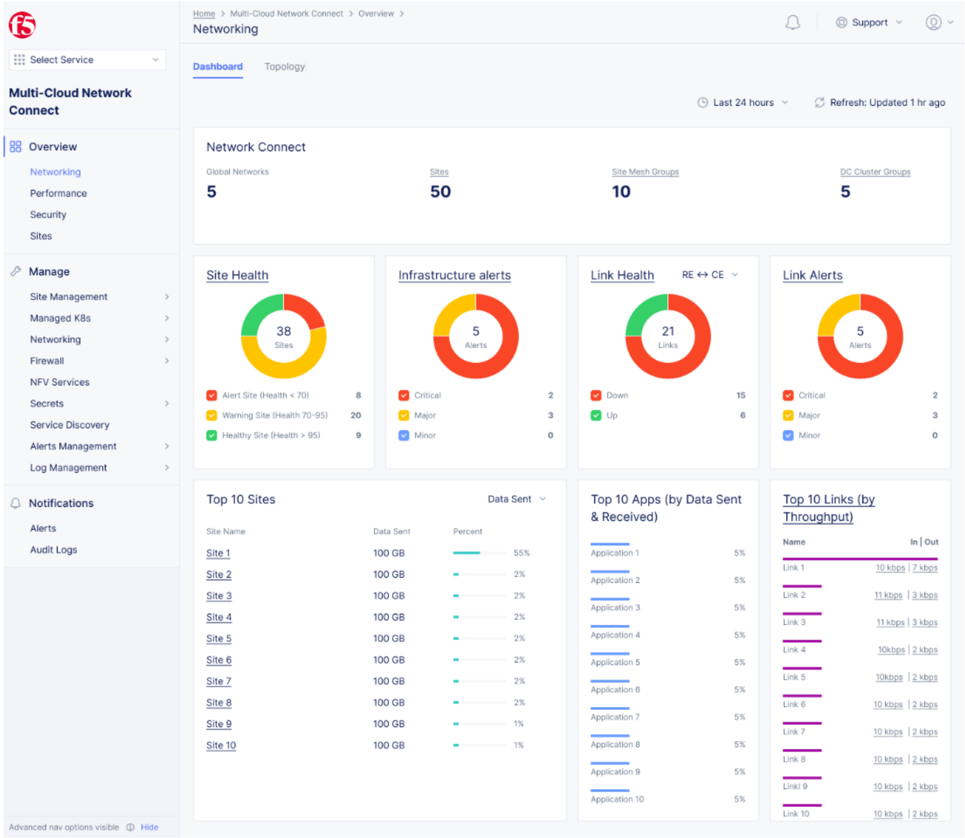

Network Connect

Global Networks	Sites	Site Mesh Groups	DC Cluster Groups
5	50	10	5

Site Health

Alert Site (Health < 70)	8
Warning Site (Health 70-95)	20
Healthy Site (Health > 95)	9

Infrastructure alerts

Critical	2
Major	3
Minor	0

Link Health RE ↔ CE

Down	15
Up	6

Link Alerts

Critical	2
Major	3
Minor	0

Top 10 Sites Data Sent

Site Name	Data Sent	Percent
Site 1	100 GB	55%
Site 2	100 GB	2%
Site 3	100 GB	2%
Site 4	100 GB	2%
Site 5	100 GB	2%
Site 6	100 GB	2%
Site 7	100 GB	2%
Site 8	100 GB	2%
Site 9	100 GB	1%
Site 10	100 GB	1%

Top 10 Apps (by Data Sent & Received)

Application 1	5%
Application 2	5%
Application 3	5%
Application 4	5%
Application 5	5%
Application 6	5%
Application 7	5%
Application 8	5%
Application 9	5%
Application 10	5%

Top 10 Links (by Throughput)

Name	In | Out
Link 1	10 kbps | 7 kbps
Link 2	11 kbps | 3 kbps
Link 3	11 kbps | 3 kbps
Link 4	10kbps | 2 kbps
Link 5	10kbps | 2 kbps
Link 6	10 kbps | 2 kbps
Link 7	10 kbps | 2 kbps
Link 8	10 kbps | 2 kbps
Linkl 9	10 kbps | 2 kbps
Link 10	10 kbps | 2 kbps

Figure 3: New top level networking dashboard showing connectivity health, alerts, and activity of connected environments.

Application Performance and Reliability

Synthetic Monitoring

General Availability of Synthetic Monitoring

F5 Distributed Cloud Synthetic Monitoring becomes generally available as part of Observability. Synthetic Monitoring is an easy-to-use service that significantly reduces mean time to resolution of application issues through uptime, performance, and health analytics. For more information on how to use it, see the Synthetic Monitoring User Guide.

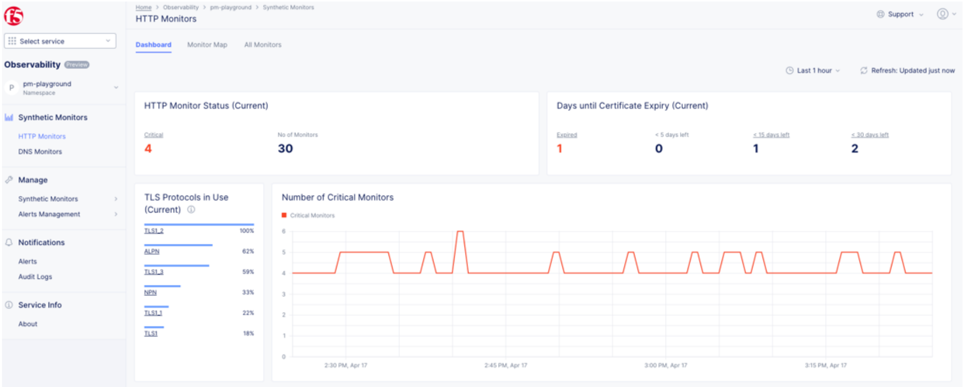

HTTP Monitor Status (Current)

Critical	No of Monitors
4	30

Days until Certificate Expiry (Current)

Expired	< 5 days left	≤ 15 days left	≤ 30 days left
1	0	1	2

TLS Protocols in Use (Current)

TLS1_2	100%
ALPN	62%
TLS1_3	59%
NPN	33%
TLS1_1	22%
TLS1	18%

Figure 4: HTTP(s) monitors from F5 Distributed Cloud Synthetic Monitoring simulate user HTTP(s) requests to gauge the health and performance of customer endpoints worldwide.

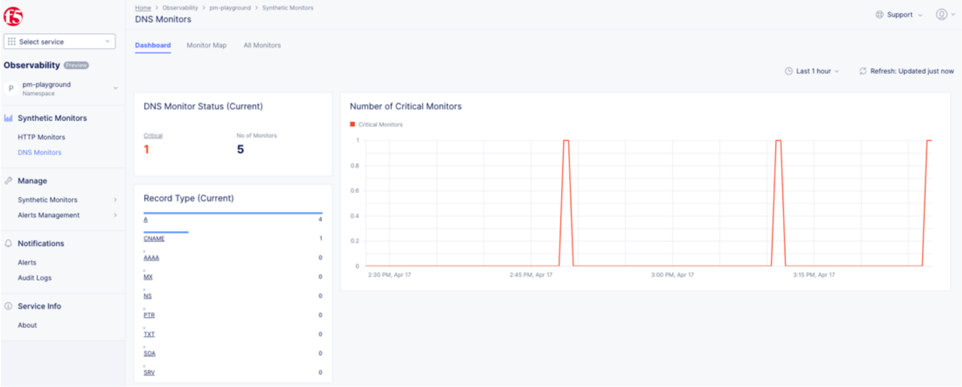

DNS Monitor Status (Current)

Critical	No of Monitors
1	5

Record Type (Current)

A	4
CNAME	1
AAAA	0
MX	0
NS	0
PTR	0
TXT	0
SOA	0
SRV	0

Figure 5: DNS monitors from Synthetic Monitoring replicate user DNS requests to gauge the health and performance of customer endpoints worldwide.

Please see the full Distributed Cloud Changelog for additional information, including more new enhancements plus known issues and caveats. We hope you find the information contained in these release notes useful. If you have any feedback, please email CS_DistributedCloudTeam@f5.com.