Report of Failure Mode

Date of Failure Incident: February 29, 2024
Failure Event Description: Failure of ingress TCP/IP to datacenter; east-west traffic unaffected; outbound email unaffected
Primary Event: Planned upgrade event on network providers equipment (NPE)

Facts Prior to and During the Primary Event

These are known facts about the network infrastructure up to and during the time of the primary event triggering the failure.

FACT 1: On February 28, 2024, at 11PM MST, a scheduled firmware upgrade to ISP providers gateway/router (NPE) would occur.

FACT 2: The upgrade would cause a necessary outage of both primary and secondary links facilitated by the NPE device.

FACT 3: There is no backup/alternate NPE device installed.

FACT 3: The duration of service outage would be approximately 40 minutes.

FACT 4: The firmware upgrade would contain changes to L3 routing protocols (BGP, et. Al.) supporting changes to the vendor network.

FACT 5: PSFA leased IP address block is virtual and programmed to the NPE gateway device that was upgraded.

FACT 6: No documented changes occurred to CPE configurations directly connected to the vendors NPE.

FACT 7: TCP/IP ingress traffic is L2 NAT to internal network subnets.

FACT 8: TCP/IP service hosts (web servers, etc.) are defined in a NAT routing table hosted by the Firewall High-Availability stack.

FACT 9: Downstream interconnect from NPE to CPE is layer 2.

Reported Observations Post-Primary Event

These are repeatable and verifiable discoveries made during initial diagnostic routines that help describe what is going on during the primary event.

REPORT 1: East-West traffic through datacenter unaffected and nominal.

REPORT 2: Public IP statically assigned on directly connected devices consistently responded to ICMP (ping) even after power cycling NPE.

REPORT 3: Internal switch gear reported their public IP as the VIP (virtual IP) assigned to L2TP/IPSec port used for VPN access.

REPORT 4: ICMP responses occur only on some NAT services (web servers, et. al.).

REPORT 5: Traceroute reports show TCP/IP ingress traffic stopping at the next-connected device from the provider's network (ZAYO) for IPs failing to respond to ICMP.

Report of Failure Mode

Date of Failure Incident: February 29, 2024
Failure Event Description: Failure of ingress TCP/IP to datacenter; east-west traffic unaffected; outbound email unaffected
Primary Event: Planned upgrade event on network providers equipment (NPE)

REPORT 6: The next-connected device from providers network (ZAYO) is the NPE installed at PSFA datacenter.

Ancillary Observations

Additional items observed where a determination cannot be made as to whether they are relevant to the primary event. Usually, these are observations that cannot be repeated reliably or are transient by some other mode.

REPORT 1: When NPE was power cycled, a few devices behind NAT responding to ICMP would stop responding, while others would then begin to respond.

REPORT 2: Download (ingress) speed test over wired LAN nodes is severely impacted while upload (egress) is nominal.

REPORT 3: After VPN connection was established, TC/IP traffic over the tunnel would not remain stable while in use accessing services.

REPORT 4: One IP at .234 seems to be associated with duplicate DNS records and was determined to be a result of internal DSD lab systems – not production – and deemed not relevant to present event.

Report Description

At 7:12AM MST on Thursday, February 29, 2024, a report was made to the Digital Strategies Division (DSD) that external network services were not responding; namely, VPN and webservices hosted by the PSFA on-premises cloud infrastructure (PSFA*Chat*, PSFA*Connect*, nmpsfa.org).

Following the report, DSD analysts began a usual routine of verifying services which includes using an outside network to connect to internal services in effort to replicate the reported failures. This was performed and the failures reported were verified within 60 minutes.

DSD Protocols

DSD maintains a protocol for business interruptions caused by service outages, the criticality of which is determined by the scope and scale of the interruption. This interruption rated a “critical” (RED) outage affecting the whole of the enterprise to some degree or another, with no known mitigation at the time of the report.

Determination of the event as “Critical” was determined by total outage from external services, and partial outage regarding internal services such as outbound email, a critical application; the failure was system-wide and work could not be produced including a negative downstream business process failure and at the time of the report, there was no known mitigation as the failure had never occurred before.

Report of Failure Mode

Date of Failure Incident: February 29, 2024
Failure Event Description: Failure of ingress TCP/IP to datacenter; east-west traffic unaffected; outbound email unaffected
Primary Event: Planned upgrade event on network providers equipment (NPE)

DSD Communications Response

The general SLA response requires DSD to acknowledge the reporting parties report; this was done as the reporting party personally telephoned the CTO/CIO about the failure.

The qualitative assessment phase of the response protocol began which requires DSD analysts to notify external security and network maintenance and engineering partner (INS, LLC) to assist in diagnosis and remediation. Concurrently, DSD notified the ISP vendor to report the interruption of nominal network services following their scheduled upgrade to the NPE plant equipment.

As part of this phase of investigation, a communication timeline is established and if known, a resolution expectation is established. At the time, no known resolution could be established; however, the initial communication to the agency-at-large was attempted through the only remaining service available to use, cell phone push-notification. At 7:58, an email was distributed to the staff at large about the outage and communication expectations. This email was likely only received by staff at ABQ prior to our determination that inbound email was also affected. At 11:37 AM, upon discovering inbound email was impacted, DSD distributed to 37 cell phones a push-notification announcing the outage and expectations.

DSD continued to push-notify cell phones regarding updates until email services were reestablished.

Response Team Actions

In accordance with SLA protocols, DSD first works to restore services to some level of useability at the expense of efficiency by working to lower criticality with respect to operations, first to an “orange” level and then “yellow”. “Yellow” condition describes a medium outage where a work-around is known to exist; however, some access to service or performance may be inefficient.

An “Orange” condition restores some services to one or more divisions on a regional or local scale, but a known solution still is not known.

To work on lowering criticality, the DSD team must first diagnose the problem and develop testing for proposed solutions, then implement the solutions developed.

At 9:16AM, the outside contractor support became engaged with the issue. The PSFA team spent 16 hours reducing criticality to a “Orange” state restoring the most-critical systems; while VPN, access to FAD, F/6 and mypsfa.org websites were still down.

The following morning of March 1, 2024, the team, after monitoring stability of the solutions put in place the night before, began implementing the solutions to restore services to all other systems affected with the sole exception of VPN.

Report of Failure Mode

Date of Failure Incident: February 29, 2024
Failure Event Description: Failure of ingress TCP/IP to datacenter; east-west traffic unaffected; outbound email unaffected
Primary Event: Planned upgrade event on network providers equipment (NPE)

By 11AM on March 1, 2024, all services other than VPN had been restored to a nominal state and verified. The condition was upgraded to "Yellow" status as VPN was still non-functional. An update on status of VPN services was communicated.

Event Response Post-Mortem Analysis

Since it was determined that the configuration of VPN played an integral part in the primary event, that service was shutdown, disabled and de-configured to allow for establishing of all other services without risk of repetition of the primary event.

A follow-on Failure Mode Analysis (FMEA) would determine root-cause, and a permanent mitigation to the primary event should it occur again.

Failure-mode Analysis began on Monday, March 4, 2024, following monitoring analysis, restoration and condition status to "Yellow". External contractors continued to monitor for stability.

Using this documentation as a starting point, Failure-Mode and Effect Analysis provides a quantitative analysis into the technicalities of the event to discover root-cause, situational conditions that enabled the failure-mode, proposal of new, or updated controls and a re-assessment of risk after controls are established, improved, or amended.

General Description of the Failure Event

Root-cause: The ISP vendor router/gateway which holds all the identity information about internal clients and services was wiped-out during a firmware upgrade and system reboot, which is normal behavior. While the system is powered on, this information is retained for 24 hours without activity, and while there is activity, the 24-hour timer is continually reset. The system nominally is never powered down and is on battery back-up for short outages. Although the system has lost upstream internet connectivity for short durations many times and for long duration twice before, it has never been power-cycled causing the clearing of all local network identity data since it was installed nearly 4 years ago.

Internal clients in the ABQ office only traverse in an east-west direction through the datacenter and never use this system except to access the general internet; in which case, the client does not care which circuit it uses and would otherwise be wholly unaffected by any problems as long as the client can reach either of the two routers, the routers will service that client for that session. This is true even if the WAN network switches dynamically from one network to another which it is designed to do.

Manual Recovery Mitigation Plan

ISSUE DESCRIPTION

When a firmware upgrade causes a condition where configurations are suspected of not having been applied, the **Meraki Security Appliances** MAY NOT be running with what otherwise would

Report of Failure Mode

Date of Failure Incident: February 29, 2924
Failure Event Description: Failure of ingress TCP/IP to datacenter; east-west traffic unaffected; outbound email unaffected
Primary Event: Planned upgrade event on network providers equipment (NPE)

appear to be the correct configurations due to the firmware upgrade having failed to apply those configurations to the presently running image.

In the Meraki dashboard, prior configurations will still appear with data in the fields; however, they may not have been applied after the firmware upgrade occurs. This will lead to the security appliances running without custom configurations which may include NAT routing configurations.

If the security device is running without NAT routing configurations, ingress traffic coming into the appliance will not route along the anticipated path and instead may attempt a route along the "gateway of last resort" (GOLR) if packets can find that route. This condition can easily trigger an **RSTP reaction** from the ingress side of the L2 port being used for uplink to a router (usually "WAN" port).

When RSTP is triggered, it will rapidly switch between port states between WAN 1 and WAN 2 (if configured) attempting to find a path inward to the service being requested by ingress packets. If a GOLR path exists, the packets will use that path to the next hop but may die there or find alternative less efficient routes to their destination.

Once the security appliance has entered RSTP, it will endlessly cycle until the condition is cured. When the security appliance is in the RSTP state, the logs will show endless switching between WAN 1 and WAN 2 ports being active. Egress packets will find their way to whatever internet ports are open but may find increased latency due to the port switching on and off. East-west traffic within the datacenter will not be noticeably impacted. Ingress traffic (mail, websites, NAS) will have the appearance of failure from the outside and services will not respond, timeout and packets discarded.

ISSUE CONDITIONS

This issue has been observed when a firmware upgrade to security appliances has occurred and custom configurations have not been re-applied post-upgrade.

MITIGATION ACTIVITY

1. Only schedule PRODUCTION firmware versions; do not apply PRODUCTION CANDIDATE releases. You may accomplish this by manually scheduling the upgrade once the notification appears in the dashboard. Reschedule as needed until the correct production version becomes available that you wish to apply. NOTE: Meraki will force the current upgrade unless it is manually scheduled. This step alone may not be enough to avoid this problem but is proactive.
2. After the image has been upgraded, verify the public IP of Meraki devices. If using an alternate WAN (WAN 2), you will see the WAN 2 public IP assigned to all Meraki switches in the dashboard. If the firmware has applied configurations and is not in RSTP, you will see instead the public IP of the physical or virtually port instead. Note: this review requires WAN 2 uplink infrastructure and is only a presumptive test. If not using WAN 2, review the Security Appliance logs.

Report of Failure Mode

Date of Failure Incident: February 29, 2924
Failure Event Description: Failure of ingress TCP/IP to datacenter; east-west traffic unaffected; outbound email unaffected
Primary Event: Planned upgrade event on network providers equipment (NPE)

3. Review the security appliance logs and look for RSTP entries continually updating. Some RSTP upon start-up may be normal; however, if they should not be continuing well after the security appliance has stabilized. Constant RTSTP is an indication of configurations not having been applied especially if the configurations contain NAT routing.
4. Verify that websites and ICMP are consistently responsive.
5. Remediating may require physically disconnecting WAN 2 from the security appliances. This will help stop RSTP from responding and bring the Meraki plant to the primary WAN (WAN 1) to the proper public IP. This must be done first to ensure configurations synch with the Meraki cloud.
6. Wait for switches to acquire the correct public IP.
7. Navigate to the NAT routing table and hit the SAVE button at the bottom of the page to recommit the NAT routing table to the running image and the Meraki dashboard. Alternatively, each NAT route may be temporarily modified and using the pop-up SAVE button, that will usually update just the record that was changed. It is critical that each NAT routing record, or preferably the entire table be re-saved.
8. Repeat this procedure for each custom configuration dialog page (VPN, VIP, etc.)
9. If disconnected in the prior step, reconnect the WAN 2 port. Meraki switches should remain on the primary WAN 1 port.
10. Verify ICMP and websites, mail are working over HTTP/HTTPS.
11. Verify VPN access.
12. Verify in logs that RSTP is not running or has not been re-triggered.
13. Note that Meraki cloud configurations may require several minutes to complete synchronization with the Meraki cloud.\
14. Contact Meraki support with a ticket to report the issue.

SUMMARY

In the PSFA environment, ingress is Layer-2 downstream from the ISP vendor gateway/router. Ingress traffic, once on layer-2 can trigger RSTP on the inbound WAN ports seeking the inward route in the absence of NAT routes. If the ingress packets are unable to find a route, RSTP will trigger switching between inward WAN ports seeking a path and may also route packets to a GOLR if one can be found. This condition presumes a security device is highly configured with specific routes for specific traffic types (VPN, etc.). Generally, a GOLR is also configured; however, if the NAT route tables do not exist, the GOLR may also not be reachable. If all VLAN routes are trunked downstream, the packets will die at the security appliance.

Report of Failure Mode

Date of Failure Incident: February 29, 2024
Failure Event Description: Failure of ingress TCP/IP to datacenter; east-west traffic unaffected; outbound email unaffected
Primary Event: Planned upgrade event on network providers equipment (NPE)

Maintenance Protocol Recommendations

Within the Firewall settings of the Meraki Security Appliance is a route table report showing route data saved in the internal database; however, these routes may become unusable, corrupt or unapplied after an upgrade. The solution after an upgrade is to rebuild route tables to be certain they are functional. It is presumed that Meraki is aware route tables may require a rebuild or the feature would be unnecessary.

Rebuilding routes after each firmware upgrade should be added to maintenance protocols to ensure resolution of any future routing issues including VPN. This "pop-up" that appears on the information icon tells the story: "*manual rebuilds are not normally required. However, if you've just changed [or had changed through firmware] your configuration . . .*"

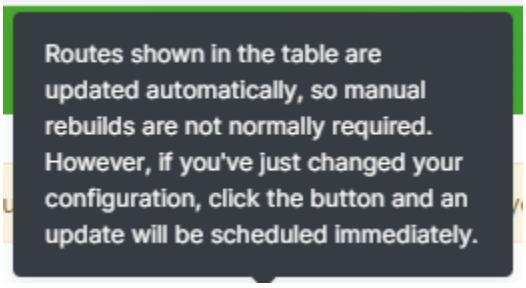

Two-factor Authentication is not currently enabled on your Meraki Dashboard account. For an extra layer of security, we recommend enabling it at your earliest convenience.

Route table Rebuild

Looking for the previ

Routes updated as of Today at 7:55 AM.

IP VERSION	SUBNET/PREFIX	NAME	VLAN	NEXT HOP	DESTINATION	TYPE	REPORTED
All	Search by subnet/prefix	Search by name	Search by VLAN ID	Search by network	Search by destination	All	Current

Stat	Version	Subnet	Name	VLAN	Next hop	Destination	Type
●	4	0.0.0.0/0	Default	—	—	WAN uplink	Default WAN Route
●	4	192.168.3.0/26	—	—	Client VPN	—	Client VPN
●	4	192.168.4.0/24	Default VLAN	4	192.168.4.1	192.168.4.1	Local VLAN
●	4	192.168.5.0/24	Vlan5	—	192.168.4.2	—	Local Static Route
●	4	192.168.11.0/24	VLAN11	—	192.168.4.2	—	Local Static Route
●	4	192.168.50.0/24	VLAN50	—	192.168.4.2	—	Local Static Route
●	4	192.168.66.0/24	Vlan66	—	192.168.4.2	—	Local Static Route
●	4	192.168.120.0/24	vlan120	—	192.168.4.2	—	Local Static Route
●	4	192.168.150.0/24	vlan150	—	192.168.4.2	—	Local Static Route
●	4	192.168.220.0/24	Vlan220	—	192.168.4.2	—	Local Static Route

10 results

Report of Failure Mode

Date of Failure Incident: February 29, 2024
Failure Event Description: Failure of ingress TCP/IP to datacenter; east-west traffic unaffected; outbound email unaffected
Primary Event: Planned upgrade event on network providers equipment (NPE)