1. information on the existing virtualisation infrastructure for virtual appliances or container-based applications
   - Concerning existing virtualization infrastructure: we are powered by VMware ESXi. We do not use any container-based applications.

2. The scope of work is to supply WAF/NLB/APG + initial deployment phase. It is also reported that this is 4 years term contract. What is to be included in the 4 years considering that supply + configuration will only take the initial part (supposedly some months only?). Is there any maintenance service to be provided?
   - We assume that there are some maintenance or subscription, or warranty fees, either from the vendor itself, or from the service provider. In such case, they have to cover the period of 4 years.

3. We understand that the required solution has to be in the form of VM appliances. Do you have on-premise servers to host the solution or you are willing to have, or do you prefer a cloud solution?
   - A: The initial requirement is to support either a virtual appliance or SaaS. In case of virtual, we will be able to host it on-premise. We will also consider a cloud solution.

4. Other WAF/NLB/APG are actually already deployed and running? Is this solution a new feature in your architecture or it will replace/coordinate with existing one?
   - There are no WAF instances deployed. NLB/APG services are currently delivered by on-premise F5 BIG-IP/APM appliances. The goal is to coordinate first, then gradually replace with the new solution.

5. Section 4.1.4: "the proposed solution must be implemented as a cluster". This holds true also in case of redundancy and HA is provided by the cloud solution itself? Or do you require in any case to have A/A A/S solution?
   - The redundancy and HA provided by the cloud solution is accepted.

6. Section 4.1.4: "[...]including assistance with the basic configuration in accordance with the suppliers’ best practices". This means it is only required to setup with basic configuration and no service for subsequent customization is required?
   - Correct. Please include a per-day rate should additional customization is requested.

7. Section 4.6: "work closely" means on-site at UPU offices? Which tasks should be executed at the headquarters? Is it possible to have an esteem of the quantity? Or will it be part of future agreements?
   - By closely we mean a dedicated person/team working on a project with us. Should physical presence is required this could be agreed later on.

8. Section 3.7: "services provided by the vendor shall be invoiced in a monthly basis": that means the total cost of the supply + configuration + licenses has to be split in 4years evenly? Also the one-time cost for initial supply?

- Correction: "services provided by the vendor shall be invoiced in a yearly basis". The one-time cost is paid at the first year, and any subsequent fees as per your question 1 should cover 4 years.

9. Section 4.1.4: "The proposed solution must integrate with existing network infrastructure". Is it possible to have the existing infrastructure design in order to better prepare the integration steps?
   - In case virtual, the proposed solution should support VMware.

10. Section 4.1.4: "The solution should provide a centralized management console for policy configuration, monitoring and reporting". The monitoring and reporting solution has to refer only to the on scope appliances provided or may monitor also other elements in the network? If yes what kind of elements and how many?
    - The solution should refer to the scope of provided appliances/services only.

11. Is it required to train UPU internal staff after deployment + basic configuration?
    - Yes.

12. Section 4.8. Travel, meal, and overnight stay costs when moving to UPU offices (if and only if agreed with UPU) will be paid by UPU. Is our understanding correct?
    - The only place that might require physical presence is located in the UPU headquarters in Bern, Switzerland.

13. Who is the owner and operator of the actual UPU infrastructure?
    - Currently it’s managed and operated internally (infrastructure and operational security teams).

14. There will be a dedicated technical interface to facilitate integration of the solution (for instance: DNS entries update, etc)?
    - Correct. All dependent services are managed by the UPU internally.

15. Section 4.1: 200 on-premises applications. Are the applications actually hosted on your private datacenters, or will they be migrated?
    - All applications are already hosted in the UPU data centers.
    -

16. Section 4.1: 30 cloud applications. Are the applications actually hosted on the cloud or they will be part of a migration?
    - All applications are already hosted in the cloud (Azure).

17. Section 4.1: total number of users. The numbers provided are an average or max total number?
    - Max total.

18. It is required to have details about traffic in the network in order to correctly dimension the initial solution. Do you have numbers about the total data rates for accessing the web applications involved?
    - We don't have any data rates collected at the moment, but would be willing to collect ones if needed.

19. Section 4.1.1: "Able to handle increasing traffic loads and scale horizontally to accommodate growing web application demands". How to handle the case if new licenses have to be provided to support the increased traffic?
    - Please provide with the pricing model that will allow us to evaluate it. For example, XX amount of calls/data is included in the package. Extra XX amount of calls/data is charged as per the following.

20. You said that you have On-Premises applications and Cloud applications: Does the cloud application needs perimeter authentication as the on-premises application. Do they need the same types of SSO?
    - Correct, cloud applications require perimeter authentication as the on-premise one. They will also need the same type of SSO. At the moment this is not the case, but will be required in the future."

21. What are the throughputs of applications on-premises and on azure?
    - We don't have a lot of apps in Azure at the moment (it's ongoing), but regarding on-premise requested graphs are attached."

22. We understand by this requirement you need to have connectivity from an Azure tenant to your private datacenter, through an IPsec or SSL VPN. Do we understand well? Or do you simply need a portal to present an application to an authenticated user without publish it directly on internet?
    - Correct, the backend server located in the private datacenter will have to be accessible via IPsec tunnel should the LB be located in the cloud (Azure)."
    -

23. We understand by this requirement that you need to route traffic to backend servers based on various and configurable criteria. Is this understanding correct? Otherwise can you give more explanations?
    - That's correct.

24. From our understanding, U2F is an open standard now superseded by CTAP2. CTAP2 is a Client To Authenticator protocol aiming to define a universal protocol to communicate between an authenticator and a software client, used notably in the context of WebAuthn. FIDO2 is a passwordless authentication standard that, when used in the context of the WebAuthn web standard for authentication, allows to perform passwordless authentication on web applications. As these do not directly refer to SSO related concepts, we would like some precisions on the expectations related to U2F/FIDO2, in which context they would be used and what is expected from the reverse proxy.

- FIDO2 (U2F) requirements were inherited from the general Access Policy Gateway (APG) specifications, however, could be omitted. OAuth 2.0 and SAML 2.0 support is mandatory."

25. How many environments (dev, uat, prod) do you want to protect with our solutions?
    - Prod and uat

26. Can these environments be hosted on the same cluster internally, or do you want to split them across different clusters?
    - Same cluster is accepted.

27. How many production applications are going to be presented on the internet and protected by the WAF capabilities?
    - The idea is to cover all publicly exposed web applications, which is ~200.
    -

28. How many API endpoints do you want to protect with our solution?

- ~50

29. On those API endpoints, how many API calls are done per days?
    - ~3.5m"

30. Could you please provide the description of the actual implementation of the BIG-IP/APM (network design, number of nodes, redundancies, services, etc.)?
    - Three, independent, virtual BIG-IP clusters in active/standby mode. Two clusters are dealing with LTM functions only (+SSL offloading). The third cluster is providing LTM (+SSL offloading) and APM functions, also acting as a IDP for on-premise ADFS.

31. Do the 3 load balancers below manage different traffic?
    - Yes, in terms of VIPs, but it’s all http/s traffic.

32. Does the new solution have to be split in the same way or is one appliance to manage all traffic ok?
    - We expect the new solution to aggregate and manage all traffic in one appliance.

33. Do all 3 LB traffic need to pass through the WAF?
    - Yes. The idea is to cover all publicly exposed web applications

34. Could you provide us with a network diagram showing the zones you have? On-prem and Azure.
    - There are two physical datacenters in Bern interconnected with Azure (Switzerland North) over VPN tunnel

35. To run the appliances or services on your premises, we can provide the necessary hardware or run the services virtually on your ESXi infrastructure. Which option would you like us to offer?
    - Virtual services are fine with us.

36. In which format does the UPU expect the invoices?
    - Monthy or yearly (depends on your payment model)

37. How is responsible for the operation of the service? Operated by UPU or Managed Service by the bidder?
    - Operated by UPU.

38. What integration stages does the UPU foresee? Test, integration or production?
    - Production and UAT

39. Are there fixed maintenance windows for the platforms mentioned above? When and how often?
    - The maintenance is performed quartely (4 times per year)

40. Is there an overview/list of the current applications?
    - Mainly SharePoint apps

41. In which Azure Regions are the 30 Cloud Apps located?
    - Switzerland North

42. Does each of these 200 web applications need its own vHost (virtual host)? e.g. www.upu.int
    - Yes

43. Does each of these 30 web applications in Azure Cloud need its own vHost (virtual host)? e.g. www.upu.int
    - Yes

44. Do the 200 web applications and 30 Azure cloud applications have to be physically connected via the same load balancer and WAF infrastructure? Or can this be separated between on-premise and cloud?
    - The idea is to terminate all publicly exposed web applications via WAF

45. Which Identity Store or Identity Provider are used for internal users and external users?
    - F5 APM instance used as a IDP for on-premise ADFS

46. Do the 10,000 external users need concurrent sessions (HTTP/S sessions)?
    - Yes

47. What is the expected growth of accessing identities, both external and internal, over the contract term?
    - N/A

48. Do the 10,000 external users need simultaneous authenticated sessions (MFA)?
    - No

49. How often per year do you expect traffic spikes and to what factor compared to the average traffic load?
    - Twice per year (2-5 weeks each) during the events (conferences)
    - The amount of users can go up to 10000 (comparing to 200-300 users during regular days), hence consider the increase of traffic up to 40-50 times.

50. What is the expected growth of the applications over the contract term?
    - Very roughly, about 3-5 apps per year.

51. What SIEM is in use at UPU, into which the NLB should be integrated for centralized logging and analysis?
    - Splunk/Sophos

52. Could you describe the use case with the UDP & FTP protocols in more detail? Would the implementation with a load balancer at Layer 3 to WAF at Layer 7 with SSL/TLS termination be sufficient?
    - L3 NLB and L7 WAF would be sufficient

53. Does UPU already have a regional (on-premise) load balancing or global server load balancing (GSLB) infrastructure in operation that can be used for the WAF and APG? If so, which one?
    - UPU uses on-premise F5 BIG-IP LTM/APM. There is no GSLB available.
    -

54. In our response, do you want a fixed price for a fixed architecture covering all you needs? Or would you rather have a pricing model allowing to understand to possible architecture variations and their impact on the final solution price?
    - A fixed architecture covering all our needs.

55. In a previous answer, you indicated that you want to protect around 50 API endpoints. In order to make this as clear as possible and propose the most accurate solution, are those 50 endpoints behind a single hostname/Virtual Server? Or are there 50 potential virtual servers that serve an API?
    - There are ~2-3 endpoints per virtual server, meaning there are ~20-25 virtual servers."

56. Could you please provide two groups of applications with their respective bandwidth usage:

    Applications that need load balancing with none/limited security needs (OWASP)
    Applications that need advanced WAF capabilities (typically machine learning)

    - All our publicly exposed web applications will have to be protected by WAF capabilities. Most of our web applications are based on SharePoint 2016/2019. Please combine three graphs below to get a cumulative bandwidth usage.

57. If possible, provide number of concurrent sessions and session setup rate, bandwidth where SSL offloading is needed per application group.
    - Please combine three graphs below to get a cumulative bandwidth and concurrent sessions usage. All our web applications benefit from SSL offloading.

58. For on premise appliances, please provide needed connectivity (1G/10G Fiber/RJ45, number of ports)?
    - Our initial requirement was a virtual appliance or SaaS. In case you still consider a physical appliance, then at least two 10 Gbps SFP+ (10GBase-SR) interfaces supporting LACP and 802.1Q encapsulation are requested.

59. Does the appliance need to encrypt towards the backend (IPsec) by itself or is this done by an external firewall?
    - Will be done by an external firewall.

60. Does the solution have to provide IdP function for the internal and/or external users?
    - No, it’s not required.

61. Would you like the migration of these solution to take place during working or non-working hours?
    - Depending on the application, but the vast majority could be migrated during working hours.

62. Do you accept the proposal of options or variants for this call for tenders?
    - Yes

63. Among the total number of applications that would use the WAF/Load balancing solution, are all applications Public/Internet facing applications ? or are there internal applications as well (non internet facing)? if yes, how many ?
    - All applications are public/Internet facing

64. The solution aims to protect applications located in different location (on-prem and cloud) and potentially on different securit zones. Could you please provide the number of zones + the number of applications per zone ? (we would need a more granular view than the total number of applications provided in your RFP).
    - We can exclude cloud (Azure) apps for the moment since the usage is very low. Concerning on-premise: we're running ~200 web servers in 2 x /24 public subnets. Most of the web servers are behind F5 NLBs providing load balancing and SSL offloading. Most of the applications are based on on-premise SharePoints 2016 or 2019.

65. Could you please provide an estimation of the traffic bandwidth for each zone ? Ie. which throughput do we need on the solution to be able to handle the UPU traffic ?
    - Find below. The sum of three would be our actual usage.

66. Could you please provide your requirement in terms of concurrent connections for each zone ? SSL Transactions per zone ?
    - find below. The sum of three would be our actual usage.

67. How many applications will be protected by the WAF ?
    - The goal is to cover all publicly/Internet exposed applications

68. For the Access policy gateway, there are solutions that also provide Remote remote access VPN. Would you also need this remote access feature ? If yes, how many concurrent VPN users do you need ? Or concurrent sessions ?
    - It's not part of the requirements, but please include it for the reference, for the 100 concurrent users.

69. About Public Cloud, do you already have a tenant with a provider ? If yes or if it is planned, please could you precise Azure, AWS or Google, and how many workloads should be protected by the WAF?
    - Azure, around 30.

70. For the Access policy gateway, which type of users would you want to authenticate ? What are your current authentication bases ?
    - We host two on-premise ADs, those users will be part of the scope. Currently we use ADFS. This might change -- there is an ongoing project to unify (or combine) both ADs.

71. Could you please provide two groups of applications with their respective bandwidth usage:

    Applications that need load balancing with none/limited security needs (OWASP)
    Applications that need advanced WAF capabilities (typically machine learning)"

    - The goal is to protect all publicly exposed web applications via WAF. Thew sum of three graphs attached would be our overall bandwidth usage that needs to handled by WAF.

72. Please provide number of concurrent sessions and session setup rate, bandwidth where SSL offloading is needed per application group.
    - See annex

73. Are multiple deployments planned. E.g. on-prem and cloud?
    - Depends on your proposal, although we'd prefer to aggregate it in one place (be it virtual appliance or SaaS)

74. For on premise appliances, please provide needed connectivity (1G/10G Fiber/RJ45, number of ports)? Our initial requirement was a virtual appliance or SaaS.
    - In case you still consider a physical appliance, then at least two 10 Gbps SFP+ (10GBase-SR) interfaces supporting LACP and 802.1Q encapsulation are requested.

75. Does the appliance need to encrypt towards the backend (IPsec) by itself or is this done by an external firewall?
    - Will be done by an external firewall.
76. How mandatory is the requirement "granular access control policies based on user roles". (A WAF can authenticate with various protocols, but the granular autorization is usually done in the backend application)
    - For example retrieving group membership during authentication

77. Please describe the high level architecture on premise and Azure, otherwise we can't do a proper sizing / BOM. Questions:

    On-premise: Do you want to achieve geo redunancy across more than one DC?

    - Attached.
    - Geo redundancy is not requested

78. Do you expect to have a dedicated load balancer cluster separated from the WAF offered? Or do you utilize existing F5 for load balancing?
    - Initially will coexist with the existing F5s, but the goal is to aggregate everything in one place.
    -
79. Is it mandatory for this solution to loadbalance other protocols than HTTP/HTTPS/FTP? (Somehow there are F5 loadbalancers mentioned that are existing)
    - Please refer to 4.1.2
    - Compatible with a wide range of application protocols, including HTTP(S), TCP, UDP and FTP"

80. If tranceivers needs to be included in the offer, please describe in detail the amount and model (which speed, fibre connectors, MM / SM).
    - Our initial requirement was a virtual appliance or SaaS.

81. Please describe, which IdP should be integrated (e.g. Entra/Azure and Active Directory)
    - On-premise ADFS

82. On-premise deployment: Do you want to achieve geo redunancy across more than one data center? If yes, can HA clusters be stretched across the data centers? How the data centers are interconnected (speed, distance/latency, layer 2 possible)
    - No, geo redundancy is not requested

83. Do you prefer hardware or virtual appliances? For VMs: Which hypervisor you are using?
    - Virtual. VMware

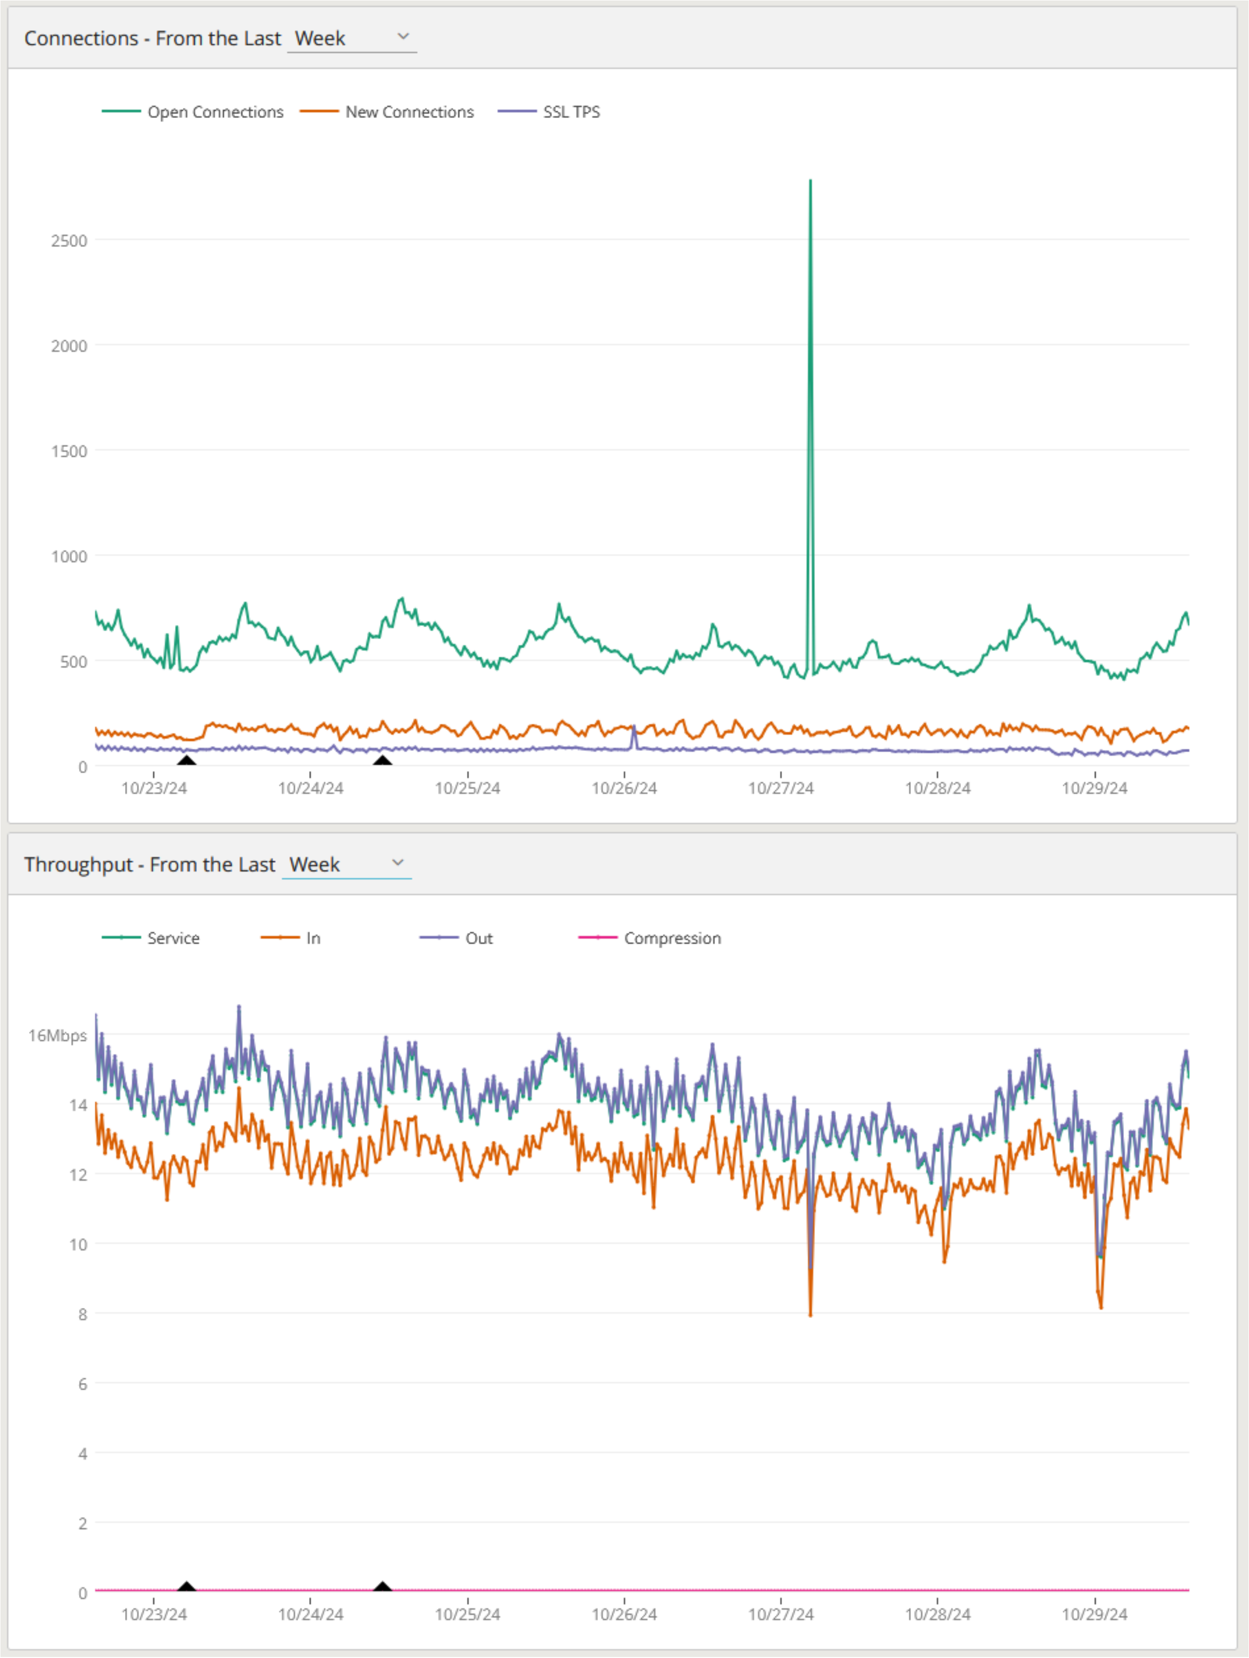

Connections - From the Last Week

Throughput - From the Last Week

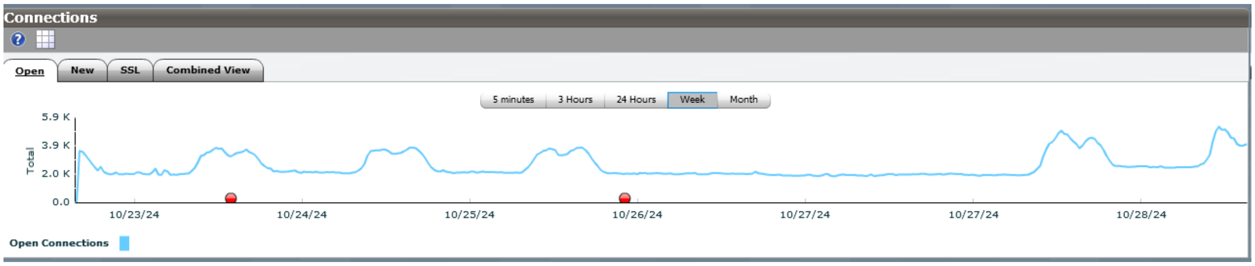

Connections

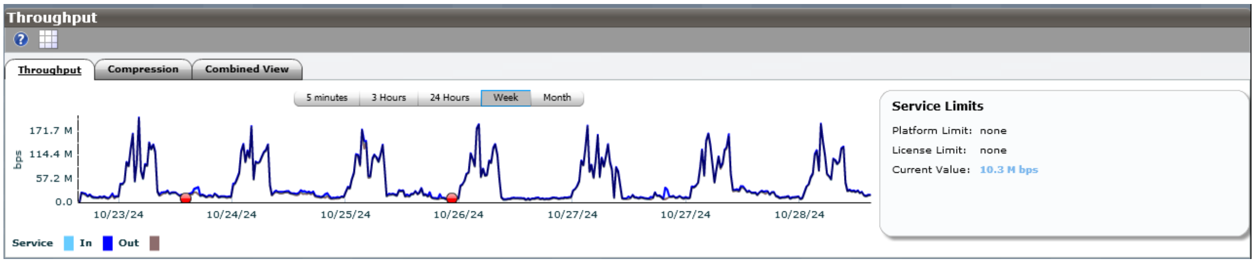

Throughput

Service Limits

Platform Limit: none

License Limit: none

Current Value: 10.3 M bps

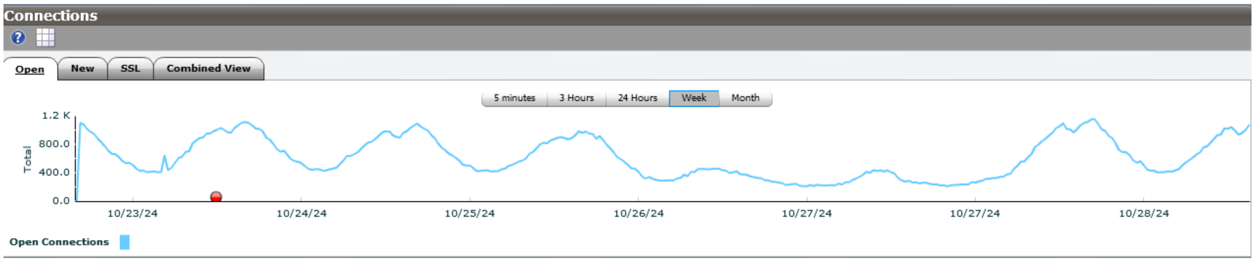

Connections

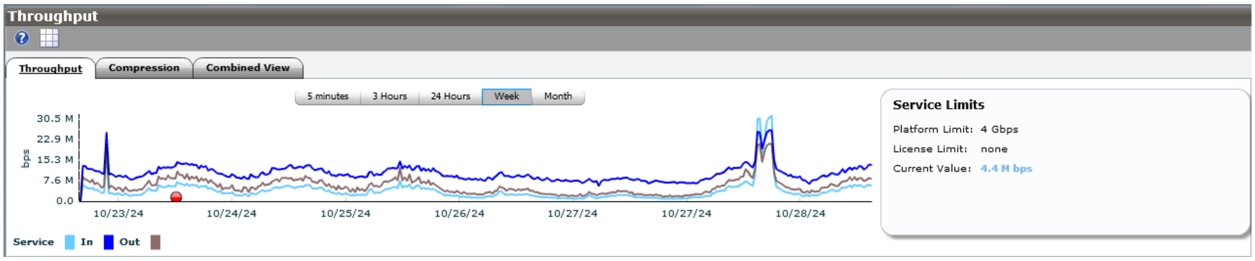

Throughput

Service Limits

Platform Limit: 4 Gbps

License Limit: none

Current Value: 4.4 M bps