Linux 性能優化實戰

39 | 案例篇：怎么缓解 DDoS 攻击带来的性能下降问题？

導讀：Earou Huang

什麼是 DDoS? (Distributed Denial of Service)

DDoS 的前身是 DoS（Denail of Service）
即拒绝服务攻击，指利用大量的合理请求，
来占用过多的目标资源，
从而使目标服务无法响应正常请求。

三種做法 x 不同的 Layer

- L3 的合法請求
- L4 的合法請求
- L7 的合法請求

耗盡
頻寬

耗盡
OS資源

耗盡
Application
資源

三種做法 X

針對 L3 設計的攻擊

- Ping flood
- Smurf attack
- Ping of death

- L3 的合法請求 ICMP
- L4 的合法請求
- L7 的合法請求

耗盡
頻寬

耗盡
OS資源

耗盡
Application
資源

三種做法 x

針對 L4 設計的攻擊

- SYN flood
- UDP flood

- L3 的合法請求
- L4 的合法請求
- L7 的合法請求

TCP/UDP

三種做法 x

針對 L7 設計的攻擊

- HTTP flood
- DNS 攻擊
- 族繁不及備載

- L3 的合法請求
- L4 的合法請求
- L7 的合法請求

HTTP等一大堆

耗盡
頻寬

耗盡
OS資源

耗盡
Application
資源

Attack Possibilities by OSI Layer

OSI Layer	Protocol Data Unit (PDU)	Layer Description	Protocols	Examples of Denial of Service Techniques at Each Level	Potential Impact of DoS Attack	Mitigation Options for Attack Type
Application Layer (7)	Data	Message and packet creation begins. DB access is on this level. End-user protocols such as FTP, SMTP, Telnet, and RAS work at this layer	Uses the Protocols FTP, HTTP, POP3, & SMTP and its device is the Gateway	PDF GET requests, HTTP GET, HTTP POST, = website forms (login, uploading photo/video, submitting feedback)	Reach resource limits of services Resource starvation	Application monitoring is the practice of monitoring software applications using a dedicated set of algorithms, technologies, and approaches to detect zero day and application layer (Layer 7 attacks). Once identified these attacks can be stopped and traced back to a specific source more easily than other types of DDoS attacks
Presentation Layer (6)	Data	Translates the data format from sender to receiver	Uses the Protocols Compression & Encryption	Malformed SSL Requests -- Inspecting SSL encryption packets is resource intensive. Attackers use SSL to tunnel HTTP attacks to target the server	The affected systems could stop accepting SSL connections or automatically restart	To mitigate, consider options like offloading the SSL from the origin infrastructure and inspecting the application traffic for signs of attacks traffic or violations of policy at an applications delivery platform (ADP). A good ADP will also ensure that your traffic is then re-encrypted and forwarded back to the origin infrastructure with unencrypted content only ever residing in protected memory on a secure bastion host
Session (5)	Data	Governs establishment, termination, and sync of session within the OS over the network (ex: when you log off and on)	Uses the Protocol Logon/Logoff	Telnet DDoS-attacker exploits a flaw in a Telnet server software running on the switch, rendering Telnet services unavailable	Prevents administrator from performing switch management functions	Check with your hardware provider to determine if there's a version update or patch to mitigate the vulnerability
Transport (4)	Segment	Ensures error-free transmission between hosts: manages transmission of messages from layers 1 through 3	Uses the Protocols TCP & UDP	SYN Flood, Smurf Attack	Reach bandwidth or connection limits of hosts or networking equipment	DDoS attack blocking, commonly referred to as blackholing, is a method typically used by ISPs to stop a DDoS attack on one of its customers. This approach to block DDoS attacks makes the site in question completely inaccessible to all traffic, both malicious attack traffic and legitimate user traffic. Black holding is typically deployed by the ISP to protect other customers on its network from the adverse effects of DDoS attacks such as slow network performance and disrupted service
Network (3)	Packet	Dedicated to routing and switching information to different networks. LANs or internetworks	Uses the Protocols IP, ICMP, ARP, & RIP and uses Routers as its device	ICMP Flooding - A Layer 3 infrastructure DDoS attack method that uses ICMP messages to overload the targeted network's bandwidth	Can affect available network bandwidth and impose extra load on the firewall	Rate-limit ICMP traffic and prevent the attack from impacting bandwidth and firewall performance
Data Link (2)	Frame	Establishes, maintains, and decides how the transfer is accomplished over the physical layer	Uses the Protocols 802.3 & 802.5 and it's devices are NICs, switches bridges & WAPs	MAC flooding -- inundates the network switch with data packets	Disrupts the usual sender to recipient flow of data -- blasting across all ports	Many advances switches can be configured to limit the number of MAC addresses that can be learned on ports connected to end stations; allow discovered MAC addresses to be authenticated against an authentication, authorization and accounting (AAA) server and subsequently filtered
Physical (1)	Bits	Includes, but not limited to cables, jacks, and hubs	Uses the Protocols 100Base-T & 1000 Base-X and uses Hubs, patch panels, & RJ45 Jacks as devices	Physical destruction, obstruction, manipulation, or malfunction of physical assets	Physical assets will become unresponsive and may need to be repaired to increase availability	Practice defense in-depth tactics, use access controls, accountability, and auditing to track and control physical assets

https://us-cert.cisa.gov/sites/default/files/publications/DDoS%20Quick%20Guide.pdf

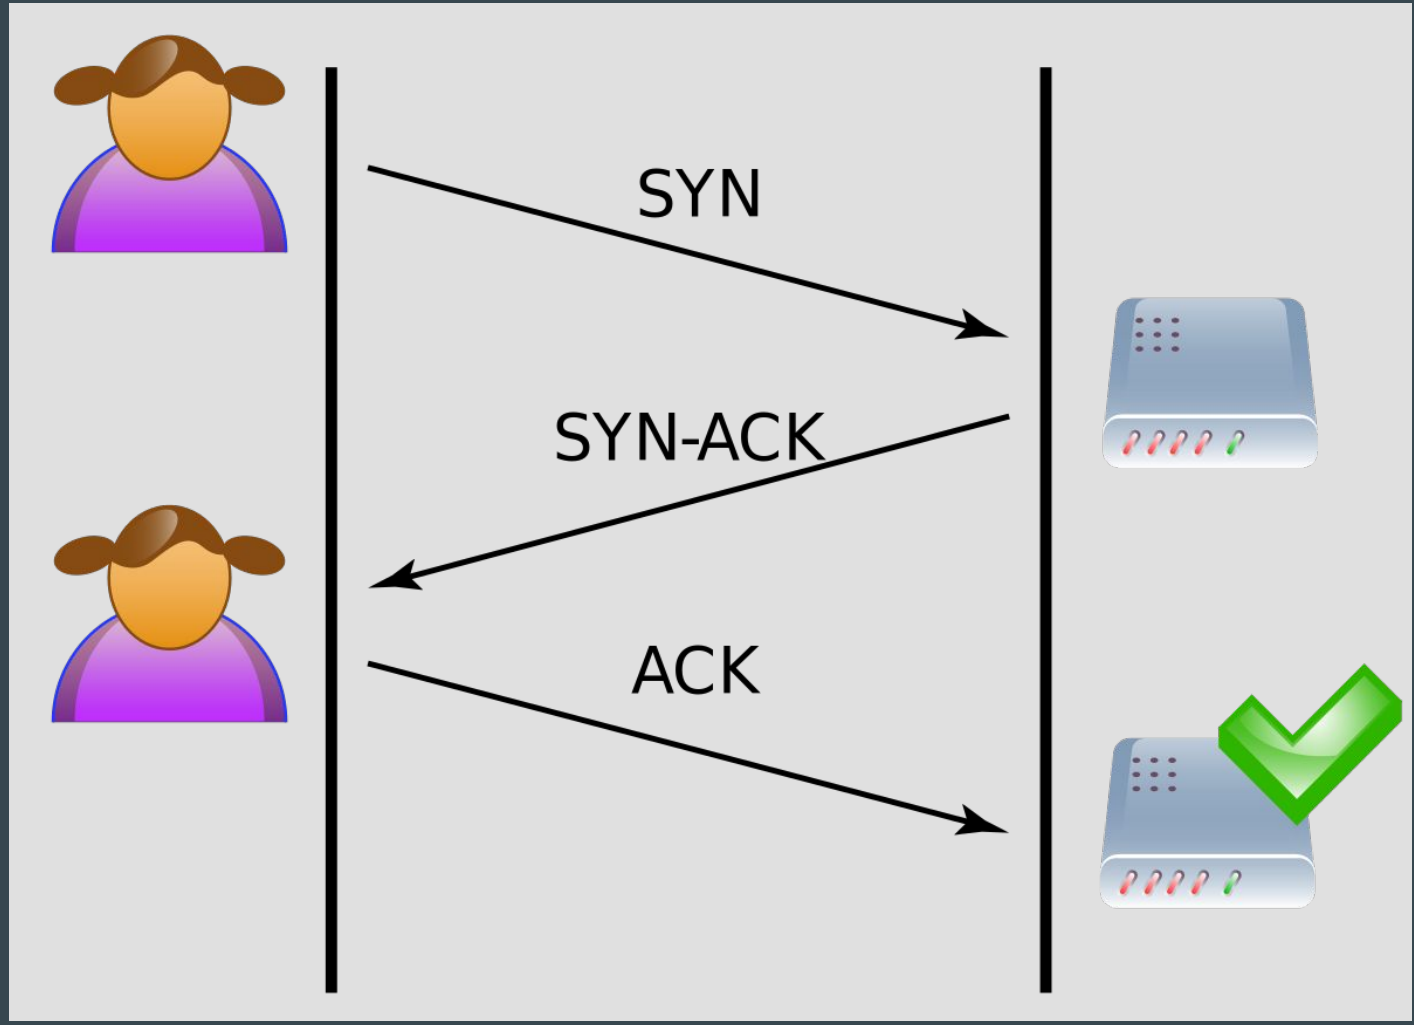
SYN
SYN-ACK
ACK

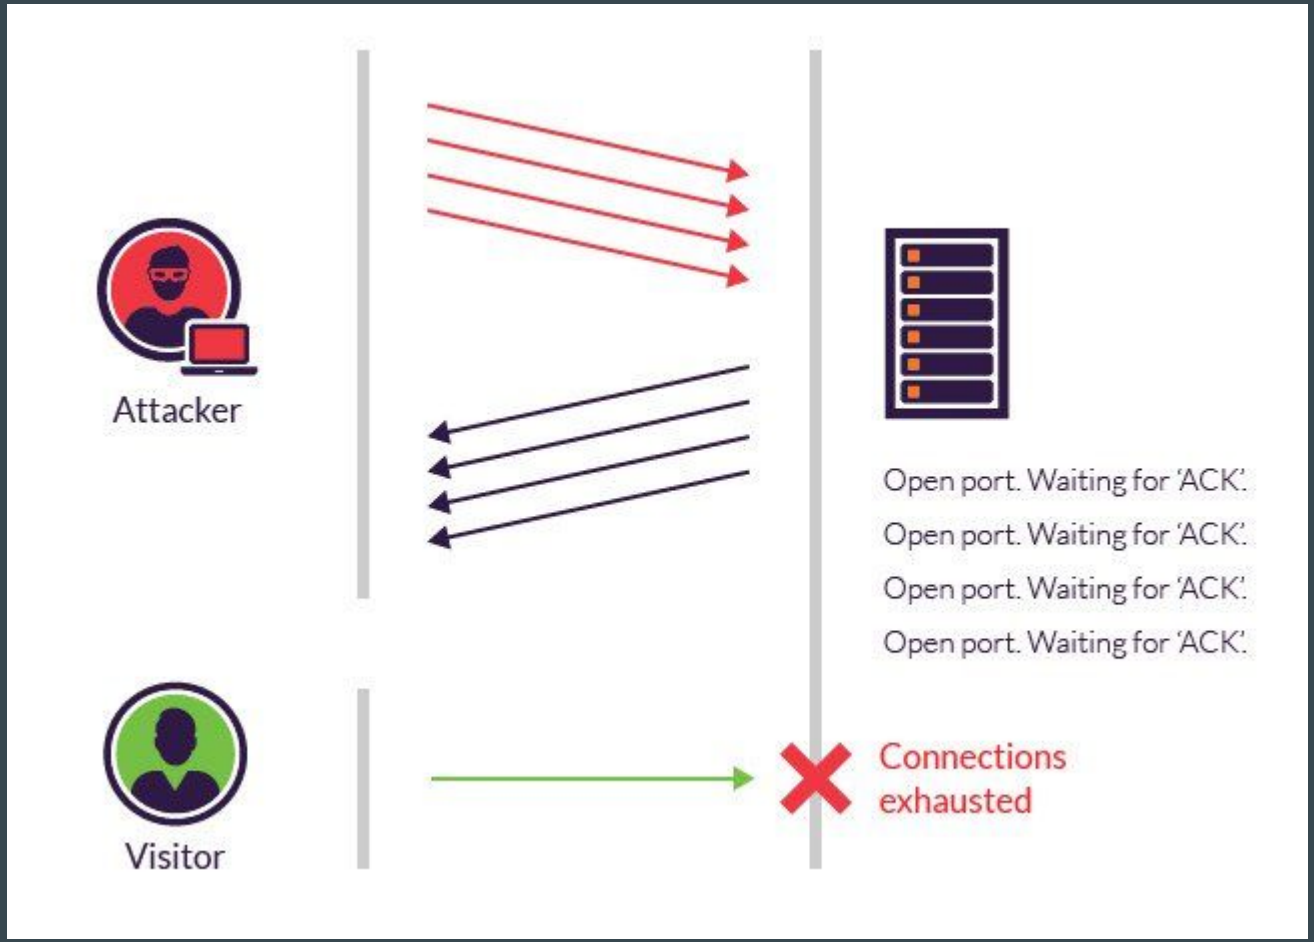
Attacker
Visitor
Open port. Waiting for 'ACK'.
Open port. Waiting for 'ACK'.
Open port. Waiting for 'ACK'.
Open port. Waiting for 'ACK'.
Connections exhausted

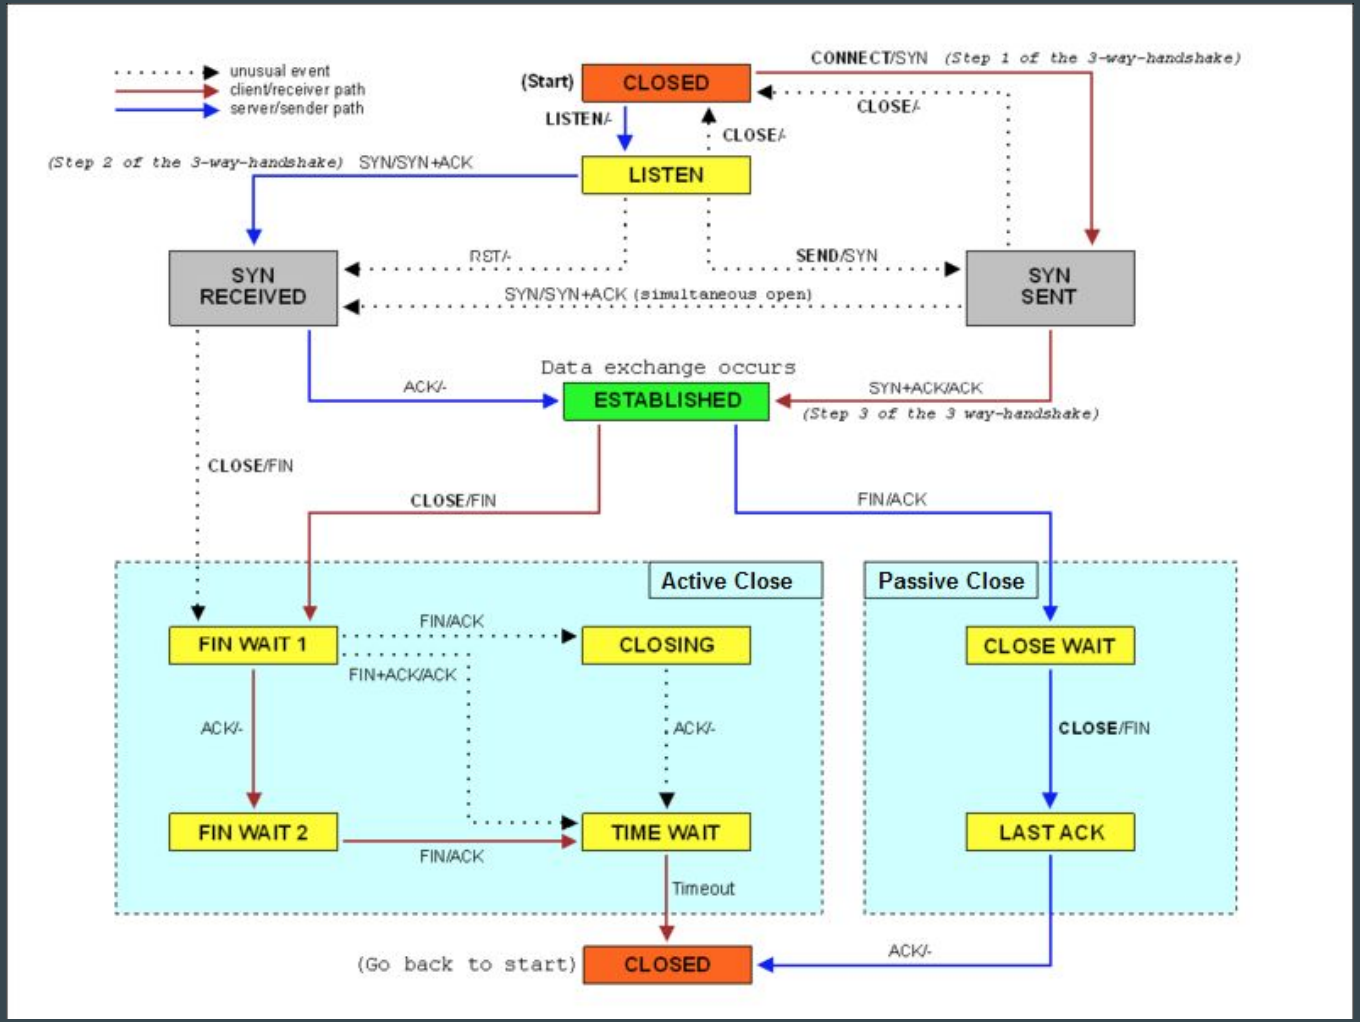
unusual event
client/receiver path
server/sender path
(Start)
CLOSED
CONNECT/SYN (Step 1 of the 3-way-handshake)
CLOSE/-
LISTEN/-
CLOSE/-
(Step 2 of the 3-way-handshake) SYN/SYN+ACK
LISTEN
RST/-
SEND/SYN
SYN RECEIVED
SYN/SYN+ACK (simultaneous open)
SYN SENT
Data exchange occurs
ACK/-
ESTABLISHED
SYN+ACK/ACK
(Step 3 of the 3 way-handshake)
CLOSE/FIN
CLOSE/FIN
FIN/ACK
Active Close
Passive Close
FIN WAIT 1
FIN/ACK
CLOSING
CLOSE WAIT
FIN+ACK/ACK
ACK/-
ACK/-
CLOSE/FIN
FIN WAIT 2
TIME WAIT
LAST ACK
FIN/ACK
Timeout
(Go back to start)
CLOSED
ACK/-

Lab

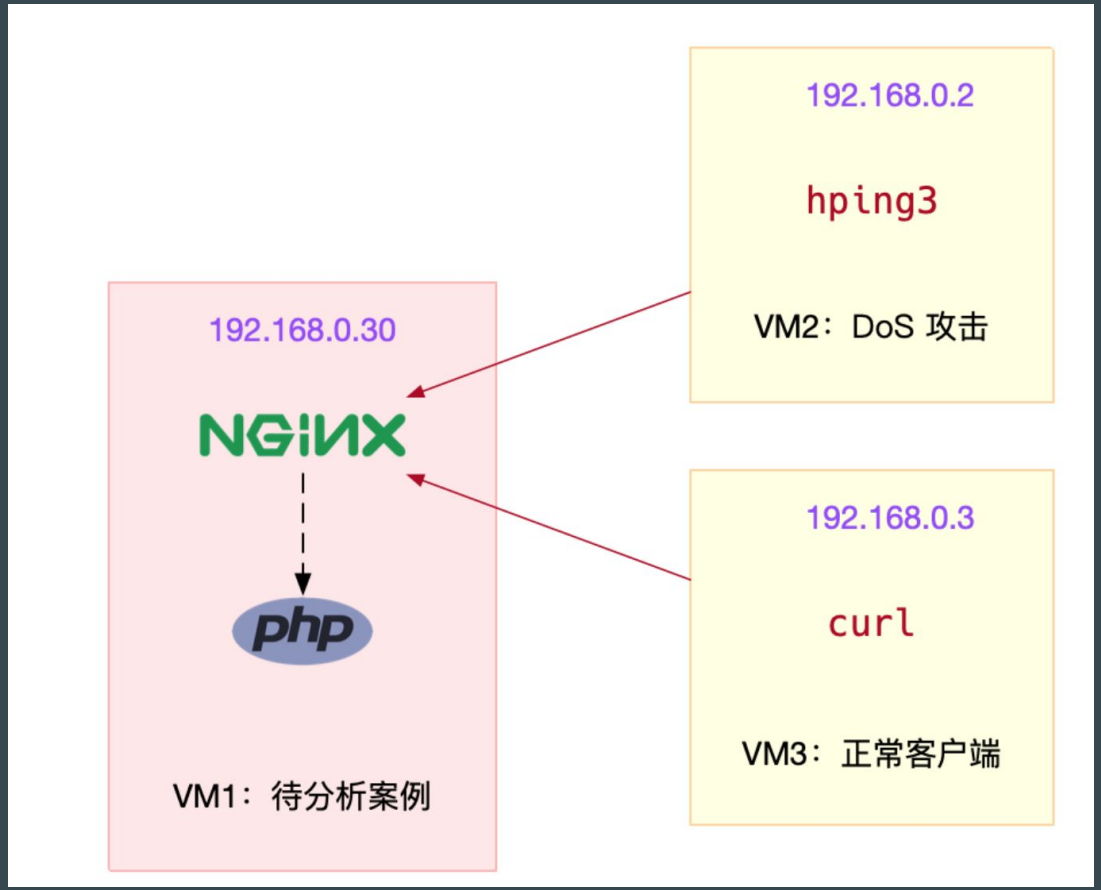
192.168.0.2
hping3
VM2：DoS 攻击
192.168.0.30
NGINX
php
VM1：待分析案例
192.168.0.3
curl
VM3：正常客户端

Takeaway 討論題

- 如果有人問, 你所維運的服務能**抵抗 DDoS 的能力**到什麼程度, 如何回答？
- 如何分辨是 **DDoS** 還是 **真正大量的使用者**?
  - 如何適當監控？(L3, L4, L7)

Addendum

Topic	Metric	Description
AWS Shield Advanced	DDoSDetected	Indicates a DDoS event for a specific Amazon Resource Name (ARN).
AWS Shield Advanced	DDoSAttackBitsPerSecond	The number of bytes observed during a DDoS event for a specific Amazon Resource Name (ARN). This metric is only available for layer 3/4 DDoS events.
AWS Shield Advanced	DDoSAttackPacketsPerSecond	The number of packets observed during a DDoS event for a specific Amazon Resource Name (ARN). This metric is only available for layer 3/4 DDoS events.
AWS Shield Advanced	DDoSAttackRequestsPerSecond	The number of requests observed during a DDoS event for a specific Amazon Resource Name (ARN). This metric is only available for layer 7 DDoS events and is only reported for the most significant layer 7 events.
AWS WAF	AllowedRequests	The number of allowed web requests.
AWS WAF	BlockedRequests	The number of blocked web requests.
AWS WAF	CountedRequests	The number of counted web requests.
Amazon CloudFront	Requests	The number of HTTP/S requests
Amazon CloudFront	TotalErrorRate	The percentage of all requests for which the HTTP status code is 4xx or 5xx.
Amazon Route 53	HealthCheckStatus	The status of the health check endpoint.

Topic	Metric	Description
ALB	ActiveConnectionCount	The total number of concurrent TCP connections that are active from clients to the load balancer, and from the load balancer to targets.
ALB	ConsumedLCUs	The number of load balancer capacity units (LCU) used by your load balancer.
ALB	HTTPCode_ELB_4XX_Count HTTPCode_ELB_5XX_Count	The number of HTTP 4xx or 5xx client error codes generated by the load balancer.
ALB	NewConnectionCount	The total number of new TCP connections established from clients to the load balancer, and from the load balancer to targets.
ALB	ProcessedBytes	The total number of bytes processed by the load balancer.
ALB	RejectedConnectionCount	The number of connections that were rejected because the load balancer had reached its maximum number of connections.
ALB	RequestCount	The number of requests that were processed.
ALB	TargetConnectionErrorCount	The number of connections that were not successfully established between the load balancer and the target.
ALB	TargetResponseTime	The time elapsed, in seconds, after the request left the load balancer until a response from the target was received.
ALB	UnHealthyHostCount	The number of targets that are considered unhealthy.

Topic	Metric	Description
NLB	NewFlowCount	The total number of new TCP flows (or connections) established from clients to targets in the time period.
NLB	ProcessedBytes	The total number of bytes processed by the load balancer, including TCP/IP headers.
Auto Scaling	GroupMaxSize	The maximum size of the Auto Scaling group.
Amazon EC2	CPUUtilization	The percentage of allocated EC2 compute units that are currently in use.
Amazon EC2	NetworkIn	The number of bytes received by the instance on all network interfaces.
NLB	ActiveFlowCount	The total number of concurrent TCP flows (or connections) from clients to targets.
NLB	ConsumedLCUs	The number of load balancer capacity units (LCU) used by your load balancer.

如何判斷何時告警？

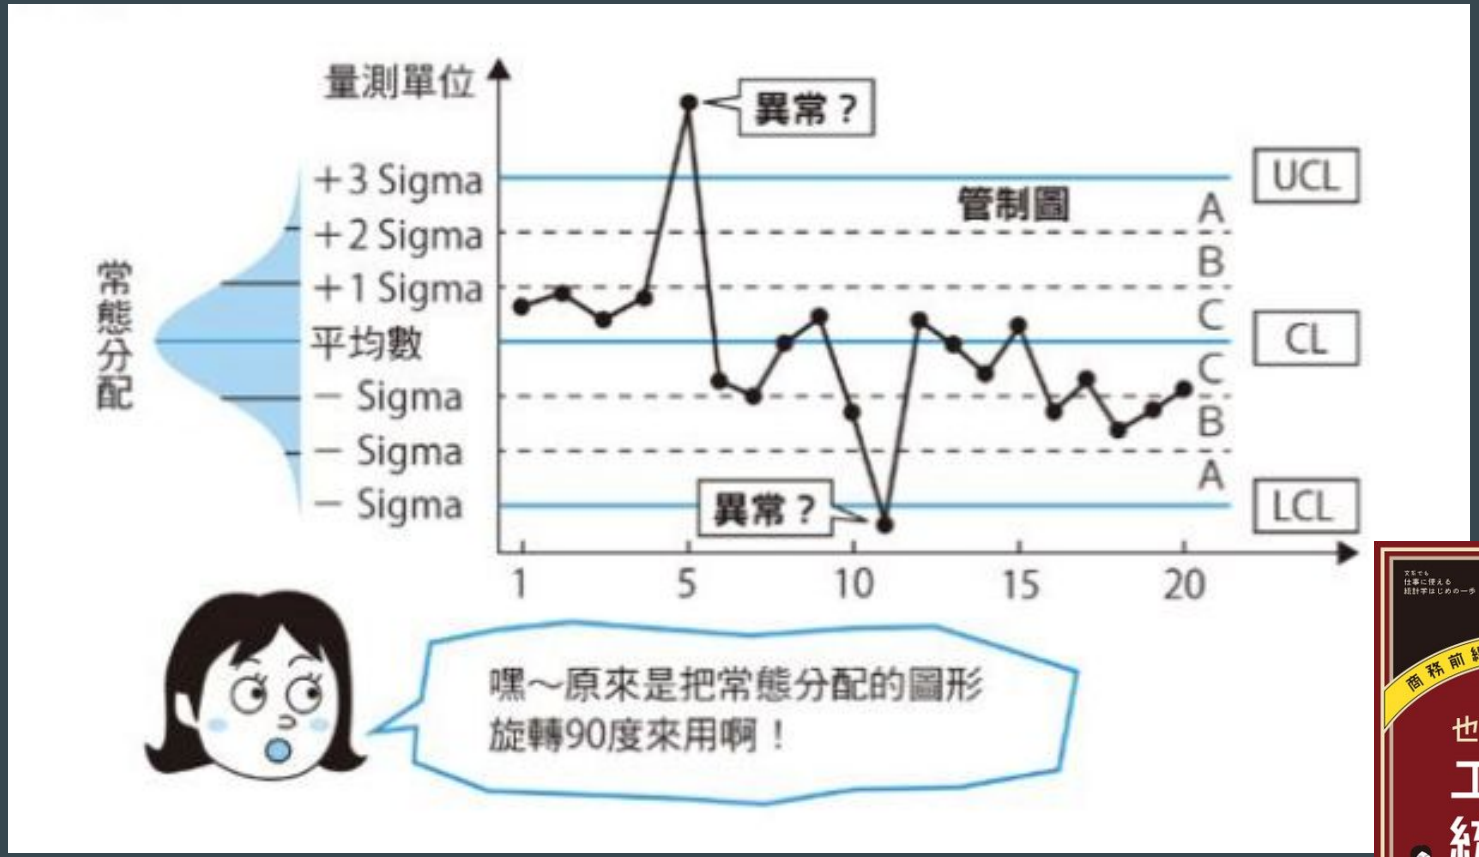

截圖自：**文科生也看得懂的工作用統計學**

Amazon CloudWatch Anomaly Detection