

在vEdge上实施出口路径时的OMP路径选择问题

目录

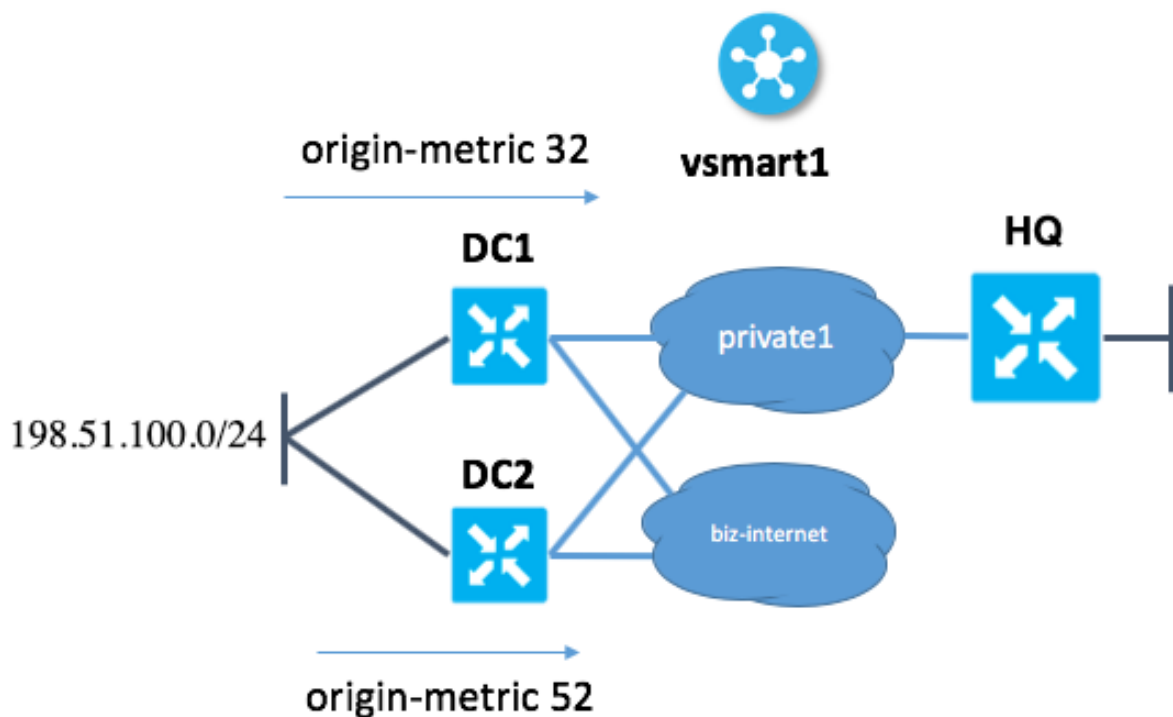
- [简介](#)
- [拓扑](#)
- [配置](#)
- [问题](#)
- [解决方案](#)

简介

本文档介绍在vEdge设备上实施重叠管理协议(OMP)路径选择而不是在vSmart控制器上实施冗余设计时出现的问题，该设计会导致不需要的结果，并在链路出现故障时导致远程站点的可达性丢失，即使备用路径可用。此问题也称为“vSmart不考虑远程vEdge上的TLOC状态”。

拓扑

为了更好地了解问题，以下是描述设置的简单拓扑图：



配置

在这里，您可以找到配置的简要说明。

- 站点DC1具有TLOC颜色“private1”和“biz-internet”
- 站点DC2具有TLOC颜色“private1”和“biz-internet”
- 站点HQ仅具有TLOC颜色“private1”
- 在DC1和DC2中，两种颜色都用于控制与vSmart的连接

两个DC站点 (DC1和DC2) 通告同一网络198.51.100.0/24。

在每个站点中，vEdge通过某种动态路由协议(例如边界网关协议(BGP))从DC获取路由器。

每个DC站点使用不同的度量标记前缀：

站点DC1 vEdge设置原始度量32

站点DC2 vEdge设置原点 — 度量52

主机名	站点ID	system-ip
DC1	21	10.100.0.21
DC2	41	10.100.0.41
总部	100	10.100.0.100
vSmart	100	10.100.0.20

问题

正常运行时：

1.vSmart从DC1和DC2接收198.51.100.0/24。

```
vsmart1# show omp routes 198.51.100.0/24
```

Code:

```
C -> chosen
I -> installed
Red -> redistributed
Rej -> rejected
L -> looped
R -> resolved
S -> stale
Ext -> extranet
Inv -> invalid
Stg -> staged
U -> TLOC unresolved
```

VPN	PREFIX	FROM PEER	PATH	STATUS	ATTRIBUTE	TLOC	IP
COLOR	ENCAP	PREFERENCE	ID LABEL		TYPE		
3	198.51.100.0/24	10.100.0.21	36 1003	C,R	installed	10.100.0.21	
biz-internet	ipsec -	<=====	METRIC 32 (PREFERRED)	10.100.0.21	49	1003	C,R
installed	10.100.0.21	private1 ipsec -	<=====	METRIC 32 (PREFERRED)	10.100.0.41	36	1003 R
installed	10.100.0.41	biz-internet ipsec -	<=====	METRIC 52	10.100.0.41	49	1003 R
installed	10.100.0.41	private1 ipsec -	<=====	METRIC 52			

2.vSmart向HQ通告目标DC1 (通过专用1和商业互联网) 的路由，因为它根据OMP路由选择标准具有最低的源度量。

```
vsmart1# show omp routes 198.51.100.0/24 vpn 3 detail
```

```
-----  
omp route entries for vpn 3 route 198.51.100.0/24  
-----
```

```
RECEIVED FROM: <===== RECEIVED FROM vEdge in DC1 in "biz-internet" color peer  
10.100.0.21 path-id 36 label 1003 status C,R loss-reason not set lost-to-peer not set lost-to-  
path-id not set Attributes: originator 10.100.0.21 type installed tloc 10.100.0.21, biz-  
internet, ipsec ultimate-tloc not set domain-id not set overlay-id 1 site-id 21 preference not  
set tag 1000030021 origin-proto eBGP origin-metric 32 as-path "65001 65001 65001" unknown-attr-  
len not set RECEIVED FROM: <===== RECEIVED FROM vEdge in DC1 in "private1" color  
peer 10.100.0.21 path-id 49 label 1003 status C,R loss-reason not set lost-to-peer not set lost-  
to-path-id not set Attributes: originator 10.100.0.21 type installed tloc 10.100.0.21, private1,  
ipsec ultimate-tloc not set domain-id not set overlay-id 1 site-id 21 preference not set tag  
1000030021 origin-proto eBGP origin-metric 32 as-path "65001 65001 65001" unknown-attr-len not  
set RECEIVED FROM: <===== RECEIVED FROM vEdge in DC2 in "biz-internet" color peer  
10.100.0.41 path-id 36 label 1003 status R loss-reason origin-metric lost-to-peer 10.100.0.21  
lost-to-path-id 49 Attributes: originator 10.100.0.41 type installed tloc 10.100.0.41, biz-  
internet, ipsec ultimate-tloc not set domain-id not set overlay-id 1 site-id 41 preference not  
set tag 1000030041 origin-proto eBGP origin-metric 52 as-path "65001 65001 65001 65001 65001"  
unknown-attr-len not set RECEIVED FROM: <===== RECEIVED FROM vEdge in DC2 in  
"private1" color peer 10.100.0.41 path-id 49 label 1003 status R loss-reason tloc-id lost-to-  
peer 10.100.0.41 lost-to-path-id 36 Attributes: originator 10.100.0.41 type installed tloc  
10.100.0.41, private1, ipsec ultimate-tloc not set domain-id not set overlay-id 1 site-id 41  
preference not set tag 1000030041 origin-proto eBGP origin-metric 52 as-path "65001 65001 65001  
65001 65001" unknown-attr-len not set ADVERTISED TO: <===== WE ADVERTISE TO HQ vEdge  
ONLY BEST ROUTES WITH METRIC 32 peer 10.100.0.100 Attributes: originator 10.100.0.21 label 1003  
path-id 4410 tloc 10.100.0.21, biz-internet, ipsec ultimate-tloc not set domain-id not set site-  
id 21 overlay-id 1 preference not set tag 1000030021 origin-proto eBGP origin-metric 32 as-path  
"65001 65001 65001" unknown-attr-len not set Attributes: originator 10.100.0.21 label 1003 path-  
id 4439 tloc 10.100.0.21, private1, ipsec ultimate-tloc not set domain-id not set site-id 21  
overlay-id 1 preference not set tag 1000030021 origin-proto eBGP origin-metric 32 as-path "65001  
65001 65001" unknown-attr-len not set
```

3. HQ vEdge将TLOC为“biz-internet”的路由标记为“Inv , U”，因为此vEdge没有TLOC biz-internet。

4. HQ vEdge将TLOC为“private1”的路由标记为“C、I、R”并安装该路由。

DC1故障场景：

1.在故障场景中，颜色为“private1”的DC1 vEdge上行链路发生故障（接口进入关闭状态），而“biz-internet”保持打开。

2.vSmart从DC1（仅可通过Biz-Internet访问）和DC2(Biz-Internet和Private1)接收198.51.100.0/24。

3.vSmart向HQ vEdge通告到DC1的路由(通过biz-internet)，因为DC1具有最低度量。

```
vsmart1# show omp routes 198.51.100.0/24 detail
```

```
-----  
omp route entries for vpn 3 route 198.51.100.0/24  
-----
```

```
RECEIVED FROM:  
peer          10.100.0.21  
path-id       36  
label         1003
```

status C,R
loss-reason not set
lost-to-peer not set
lost-to-path-id not set

Attributes:

originator 10.100.0.21
type installed
tloc 10.100.0.21, biz-internet, ipsec
ultimate-tloc not set
domain-id not set
overlay-id 1
site-id 21
preference not set
tag 1000030021
origin-proto eBGP
origin-metric 32
as-path "65001 65001 65001"
unknown-attr-len not set

RECEIVED FROM:

peer 10.100.0.41
path-id 36
label 1003
status R
loss-reason origin-metric
lost-to-peer 10.100.0.21
lost-to-path-id 36

Attributes:

originator 10.100.0.41
type installed
tloc 10.100.0.41, biz-internet, ipsec
ultimate-tloc not set
domain-id not set
overlay-id 1
site-id 41
preference not set
tag 1000030041
origin-proto eBGP
origin-metric 52
as-path "65001 65001 65001 65001 65001"
unknown-attr-len not set

RECEIVED FROM:

peer 10.100.0.41
path-id 49
label 1003
status R
loss-reason tloc-id
lost-to-peer 10.100.0.41
lost-to-path-id 36

Attributes:

originator 10.100.0.41
type installed
tloc 10.100.0.41, privatel1, ipsec
ultimate-tloc not set
domain-id not set
overlay-id 1
site-id 41
preference not set
tag 1000030041
origin-proto eBGP
origin-metric 52
as-path "65001 65001 65001 65001 65001"
unknown-attr-len not set

ADVERTISED TO:

peer 10.100.0.31

```

Attributes:
  originator      10.100.0.21
  label           1003
  path-id         5906
  tloc            10.100.0.21, biz-internet, ipsec
  ultimate-tloc   not set
  domain-id       not set
  site-id         21
  overlay-id      1
  preference      not set
  tag             1000030021
  origin-proto    eBGP
  origin-metric   32
  as-path         "65001 65001 65001"
  unknown-attr-len not set
  ADVERTISED TO:
peer 10.100.0.41

```

```

Attributes:
  originator      10.100.0.21
  label           1003
  path-id         7689
  tloc            10.100.0.21, biz-internet, ipsec
  ultimate-tloc   not set
  domain-id       not set
  site-id         21
  overlay-id      1
  preference      not set
  tag             1000030021
  origin-proto    eBGP
  origin-metric   32
  as-path         "65001 65001 65001"
  unknown-attr-len not set

```

ADVERTISED TO: <===== THIS IS WHAT WE ADVERTISE TO HQ SITE peer 10.100.0.100 Attributes:
 originator 10.100.0.21 label 1003 path-id 4410 tloc 10.100.0.21, biz-internet, ipsec ultimate-
 tloc not set domain-id not set site-id 21 overlay-id 1 preference not set tag 1000030021 origin-
 proto eBGP origin-metric 32 as-path "65001 65001 65001" unknown-attr-len not set

4. HQ vEdge将TLOC为“Inv , U”的路由标记为“Inv , U”，因为此vEdge没有TLOC biz-internet。

结果是HQ vEdge无法到达198.51.100.0/24。

解决方案

vSmart可能已向DC2发送路由（具有较低优先级的较高度量），在这种情况下，HQ vEdge仍将通过DC2使用“private1” TLOC到达目的地，DC2仍处于工作状态：

```

VEDGE-HQ-1# show bfd sessions site-id 41

```

DST PUBLIC		SITE ID		STATE		SOURCE TLOC		REMOTE TLOC		SOURCE IP			
SYSTEM IP	IP				PORT	COLOR	ENCAP	DETECT	COLOR	TX	INTERVAL(msec)	UPTIME	
10.100.0.41		41		up		private1		private1			192.168.11.1		
192.168.41.1					12406	ipsec	7			1000		12:04:02:25	0

但是，由于vSmart已选择度量较低的业务互联网路由作为最佳路径，因此，vSmart不会默认通告具有不同度量的OMP路由，因此，不允许接收vEdge设备决定采用哪条路径（并考虑可用的

TLOC)，因此，在HQ vEdge上没有通过“private1” TLOC的路由和其状态)。vSmart不考虑远程设备(在本例中为HQ vEdge)上可用的TLOC颜色，也不考虑路由的状态，因为没有这种机制来控制此情况。

这是在具有iBGP路由反射器的类似拓扑中可以看到OMP拐角情况，以及物理接口地址上的对等。

第一个解决方案选项是使用OMP中提供的类似路径的功能(RFC7911)，在vSmart上称为“[发送备份路径](#)”：

```
omp
  send-backup-paths
```

它通告所有可用路径，因此远程HQ vEdge会根据TLOC可用性选择路径。

此处的第二个解决方案选项是删除DC1和DC2 vEdge上对应前缀的路由策略操作“set metric”，然后通过vSmart控制策略执行集中式路由选择实施，如下所示：

```
policy
lists
site-list site_11
site-id 11
!
prefix-list PREFIX
ip-prefix 198.51.100.0/24
!
control-policy SET_PREF
sequence 10
match route
prefix-list PREFIX
site-id 21
!
action accept
set
preference 200
!
!
!
sequence 20
match route
prefix-list PREFIX
site-id 41
!
action accept
set
preference 100
!
!
!
default-action accept
!
apply-policy
site-list site_11
control-policy SET_PREF out
!
```

此处，site-id 11是HQ vEdge，前缀列表前缀包含您希望优先于一种或另一种TLOC颜色的前缀。由于两个OMP路由都在HQ vEdge上，因此一旦vEdge无法再到达biz-internet，它将从其OMP路由表在路由信息库(RIB)中安装通过private1的路由。