

open



USE



IMPROVE



EVANGELIZE

Crossbow: An overview

Phil Kirk
Solaris Sustaining

開
放
的
열린
مفتوح
libre
मुक्त
ಮುಕ್ತ
livre
libero
ముక్త
开放的
açık
open
nyílt
ᄒᆞᆫ
オープン
livre
ανοικτό
offen
otevřený
öppen
открытый
வெளிப்படை

Real world problems

- Large organizations doing server consolidation
- Typically hardware may have only 1-2 NIC's to share amongst many virtual machines
- If sharing the NIC becomes possible then we also need to be able to do resource control
- For certain applications there will be the need for bandwidth guarantees



Network virtualization: Crossbow

- The Crossbow architecture allows network virtualization and resource control
- NIC virtualization with VNIC's
- Stack virtualization with IP instances
- Bandwidth management
- Flow management

VNICS – Virtual NIC's

- VNIC's build on the virtualization of the MAC layer
- Allows an administrator to carve up 1Gb/s and 10Gb/s NIC's into virtual NIC's
- Each VNIC appears as a separate MAC client
- Appear to the administrator as regular NIC's
- Can be assigned to both zones and Virtual machines e.g. XVM



IP instances

- IP instances allow a zone to have it's own exclusive IP instance
- Allows totally separate network configuration for a zone, it's own routing table, IPsec configuration etc.
- VLAN's can be attached to different zones with no IP leakage between them.



Flows

- Layer 3 and above classification
- Configured based on IP address, transport, port, or diffServ field
- Can associate Bandwidth, CPU binding, Fanout, Priority
- Flows configured on a link, i.e. the MAC client
- MIB/statistics for available for flows and VNIC's so you can track usage

Administration

- Administration of VNIC's is done through `dladm(1M)`
- For flows there is a new command `flowadm`
- Configuring a zone to use IP instances is simply a matter of setting `ip-type= exclusive` using `zonecfg`

Examples

- VNIC's are managed using dladm:

```
bash-3.2# dladm create-vnic -d etherstub3 vnic10
bash-3.2# dladm show-vnic
vnic10      etherstub3      0 Mbps  2:8:20:7c:20:9d
random      0
bash-3.2#
```

- Setting bandwidth on an existing VNIC:

```
bash-3.2# dladm set-linkprop -p maxbw=100 vnic10
```


Examples cont.

- Managing flows is done using flowadm:

```
bash-3.2# flowadm add-flow -l vnic6 -a
transport=TCP,remote_port=5001 ttcp1
bash-3.2#
```

- Once we've created a flow we can then manipulate the properties for that flow:

```
bash-3.2# flowadm set-flowprop -p maxbw=20 ttcp1
bash-3.2#
```

open



USE



IMPROVE



EVANGELIZE

Thank you!

Phil Kirk

Solaris Sustaining

phil.kirk@sun.com

開
放
的
열린
مفتوح
libre
मुक्त
ಮುಕ್ತ
livre
libero
ముక్త
开放的
açık
open
nyílt
ᐅᐅᐅᐅ
𑌕𑌕𑌕𑌕
オープン
livre
ανοικτό
offen
otevřený
öppen
открытый
வெளிப்படை

“open” artwork and icons by chandan:
<http://blogs.sun.com/chandan>