

## Networking

### Network configuration - Link Aggregations

- Link aggregation is a facility for combining a number of physical network links on a system, to potentially increase the bandwidth, and provide link failover (i.e. high availability).
- For example *dladm* can be used to create a link *aggr1*, which combines links *net1*, *net2* and *net3*.
- All network configuration is then carried out on the single link *aggr1*.
- There are two types of link aggregation supported:-
  - \* Trunk (802.3ad) aggregations, which provide load spreading at MAC/IP/PORT level, but need 802.3ad network switch support, and can use only one physical switch.<sup>1</sup>
  - \* DLMP (Data Link MultiPathing) where the network links can be connected to separate switches. DLMP load spreading is at link level, and is intended primarily as the underlying link for VNICs. (Virtual Network Interfaces)
- A full explanation of the differences can be found here:-

[https://docs.oracle.com/cd/E36784\\_01/html/E37516/gmwhe.html#scrolltoc](https://docs.oracle.com/cd/E36784_01/html/E37516/gmwhe.html#scrolltoc)

---

1. Some 802.3 switches are appearing which support multiple connected switches.

## Networking

### Network configuration - Link Aggregations using 802.3ad

- An 802.3ad, or trunk, aggregation is similar to an IPMP link created for failover, but has other advantages; the main features are:-
  - \* An aggregation is treated as one device for configuration and management, and uses only one IP address.
  - \* In-bound as well as outbound load spreading is provided, so extra bandwidth is available.
  - \* Automatic failover is provided, to the remaining working links in the aggregation.
  - \* Policies can be applied to the traffic flows.
  - \* Parallel aggregations between servers can be created, without going through a switch.
- To show the available network interfaces:-

#### # *dladm show-phys*

LINK	MEDIA	STATE	SPEED	DUPLEX	DEVICE
net1	Ethernet	unknown	0	unknown	e1000g1
net2	Ethernet	unknown	0	unknown	e1000g2
net0	Ethernet	up	1000	full	e1000g0
net3	Ethernet	unknown	0	unknown	e1000g3

#### # *dladmin show-link*

LINK	CLASS	MTU	STATE	OVER
net1	phys	1500	unknown	--
net2	phys	1500	unknown	--
net0	phys	1500	up	--
net3	phys	1500	unknown	--

## Networking

### 802.3ad Trunk Aggregations - how to create

- There are conditions to consider:-
  - \* If you need to run a trunk aggregation through a switch, that switch must support 802.3ad aggregation technology.
  - \* Interfaces must be of the GLDv3 type: *xge*, *e1000g*, and *bge*. It is unlikely that other types would be supported in Solaris 11 in any case.
  - \* All interfaces in the aggregation must run at the same speed and in full duplex mode.
  - \* No IP instances must initially exist on the component interfaces.
- To set up an aggregation using *dladm*, the basics steps are:-

```
# dladm create-aggr -L active -l net1 -l net2 1
```

- \* Create an aggregation number 1 (*aggr1*) from interfaces *net1* and *net2*.
- \* *-L active* enables active LACP.

```
# ipadm create-ip aggr1
```

```
# ipadm create-addr -T static -a 130.100.1.5/16 aggr1/v4
```

- \* Brings the aggregation interface *aggr1* up with address *130.100.1.5*

```
# dladm show-aggr
```

LINK	MODE	POLICY	ADDRPOLICY	LACPACTIVITY	LACPTIMER
aggr1	trunk	L4	auto	off	short

- Use *dlstat* to display link statistics.

## Networking

### Aggregations - Modifying

- To add another interface to the aggregation:-

```
# dladm add-aggr -l net3 aggr1
```

```
# dladm show-aggr -x
```

LINK	PORT	SPEED	DUPLEX	STATE	ADDRESS	PORTSTATE
aggr1	--	1000Mb	full	up	0:14:4f:97:7d:f	--
	net1	1000Mb	full	up	0:14:4f:97:7d:f	attached
	net2	1000Mb	full	up	0:14:4f:97:7d:10	attached
	net3	1000Mb	full	up	0:14:4f:97:7d:11	attached

```
# dlstat show-aggr aggr1
```

LINK	PORT	IPKTS	RBYTES	OPKTS	OBYTES
aggr1	--	5.34M	435.78M	1.48M	190.05M
aggr1	net1	2.02M	159.87M	570.29K	73.00M
aggr1	net2	2.13M	181.61M	570.09K	72.97M
aggr1	net3	1.19M	94.31M	344.38K	44.08M

- To remove a device from the aggregation:-

```
# dladm remove-aggr -l net0 1
```

- To remove *aggr1*:-

```
# dladm delete-aggr 1
```

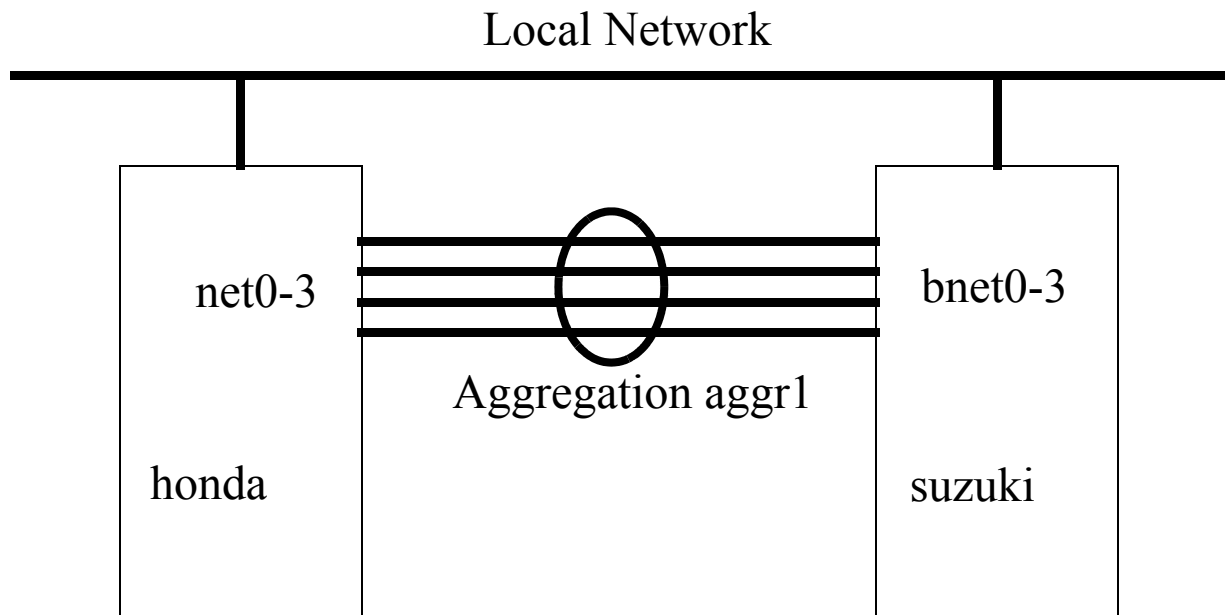
\* or..

```
# dladm delete-aggr aggr1
```

## Networking

### Aggregations - Parallel (Back-to-Back)

- A parallel, or back-to-back, aggregation can be created between servers without the need for a switch.



Systems *honda* and *suzuki* are connected via an aggregation on interfaces *net0-3*, and also to the local network through another interface.

- This is most commonly used between mirrored database servers to provide greater bandwidth between the servers themselves for database updates, etc.

## Networking

### Aggregations - Policies

- It is possible to set policies for aggregations, such as whether to use LACP (Link Aggregation Control Protocol) protocol.
- The most likely reason to change the policies would be the way the network switch handles aggregations, so we would recommend examination of the switch documentation.
- Here is an example which sets LACP mode to *active*:-

```
# dladm modify-aggr -L active aggr1
```

```
# dladm show-aggr
```

```
LINK   MODE POLICY  ADDRPOLICY LACPACTIVITY LACPTIMER
aggr1  trunk  L4       auto       active       short
```

- The default outbound load spreading (L4) is based on IP/port number hashes (hash: a smaller integer representation of a large value) but this can also be changed - see the *dladm* man page for full details.
- The switch needs to understand what mode the system is using, and support it, so LACP should be enabled on both switch and Solaris system.
- Link Aggregation implementations may differ between switches from different manufacturers.
- If there is a mis-match, you may find only a single physical link gets used; *dladm show-aggr -L* will establish this.

## Networking

### Aggregations - Zones

- Aggregations can be assigned to zones using IP instances.
- *dladm* can assign an aggregation temporarily to a zone, until the next reboot:-

```
# dladm set-linkprop -t -p zone=apache aggr1
```

- \* The *-t* option must be used.

- \* Use *zonecfg* to assign *aggr1* permanently to the *apache* zone.

## Networking

### Aggregations - DLMP

- DLMP-mode aggregations are used to support VNIC load spreading and high availability.
- VNICs are covered in detail shortly.
- To create a DLMP-mode aggregation:-

```
# dladm create-aggr -m dlmp -l net1 -l net2 -l net3 aggr1  
* -m dlmp denotes that this is a DLMP-mode  
aggregation.
```

- The links that form part of the aggregation can be connected to different switches which do not need 802.3ad support, thus eliminating the switch as a point of failure, and the need for any switch configuration.
- As VNICs are created, they are associated with an aggregation port, and if the port fails, they will be switched to another (working) port.
- If the number of VNICs exceeds the number of aggregation ports, then the ports are shared between them.
- VNICs are created as follows:-

```
# dladm create-vnic -l aggr1 vnic0
```

- More about VNICs soon.