


Network Test 4x10G

Bonding on Linux


Donato De Girolamo

Andrea Prosperini

Layout







port-channel201








Max In:35.9 Gb/s (35.9%)

Max Out:18.5 Gb/s (18.5%)


dell-test-1 / RR




dell-test-1 Load Average



dell-test-1 / MAC HASH (LACP)



diskserver



Server utilizzati

host	test
Model	Dell Inc. PowerEdge R715
CPU	2 x 16 Core AMD 6376 => 32 x AMD Opteron(tm) Processor 6376
RAM	32GB
Kernel	Linux 2.6.32-504.8.1.el6.x86_64
NIC	Broadcom Corporation NetXtreme II BCM57810 10/20 Gigabit Ethernet (bnx2x 1.78.19-0)

host	ds-xyz
System Board	Intel Corp. S5520UR
CPU	2 x Quad Core E5620 => 16 x Intel(R) Xeon(R) E5620 @ 2.40GHz
RAM	24GB
Kernel	Linux 2.6.32-504.8.1.el6.x86_64
NIC	Intel Corporation 82599 10 Gigabit Dual Port Network Connection (ixgbe)

Conclusioni

- Usando 4x10G per server possiamo ridurre di $\frac{1}{4}$ il numero dei server per l'accesso allo Storage
- Possiamo evitare di passare subito a schede/QSFP+ 40G (molto costose) visto che:
 - lo standard 40G/QSFP+ impone l'installazione di un nuovo cablaggio strutturato di tipo MPO (anch'esso molto costoso)
 - Mentre il BiDi (Cisco) non e' al momento installabile su server.
- Abbiamo margine per capire che direzione prende lo standard 40G/100G per fare poi un investimento piu'mirato:
 - sia a livello di schede
 - sia a livello di espansione dell'attuale cablaggio