

Identifier Locator Addressing

Network virtualization without encapsulation

draft-herbert-nvo3-ila-00

Tom Herbert <tom@herbertland.com>

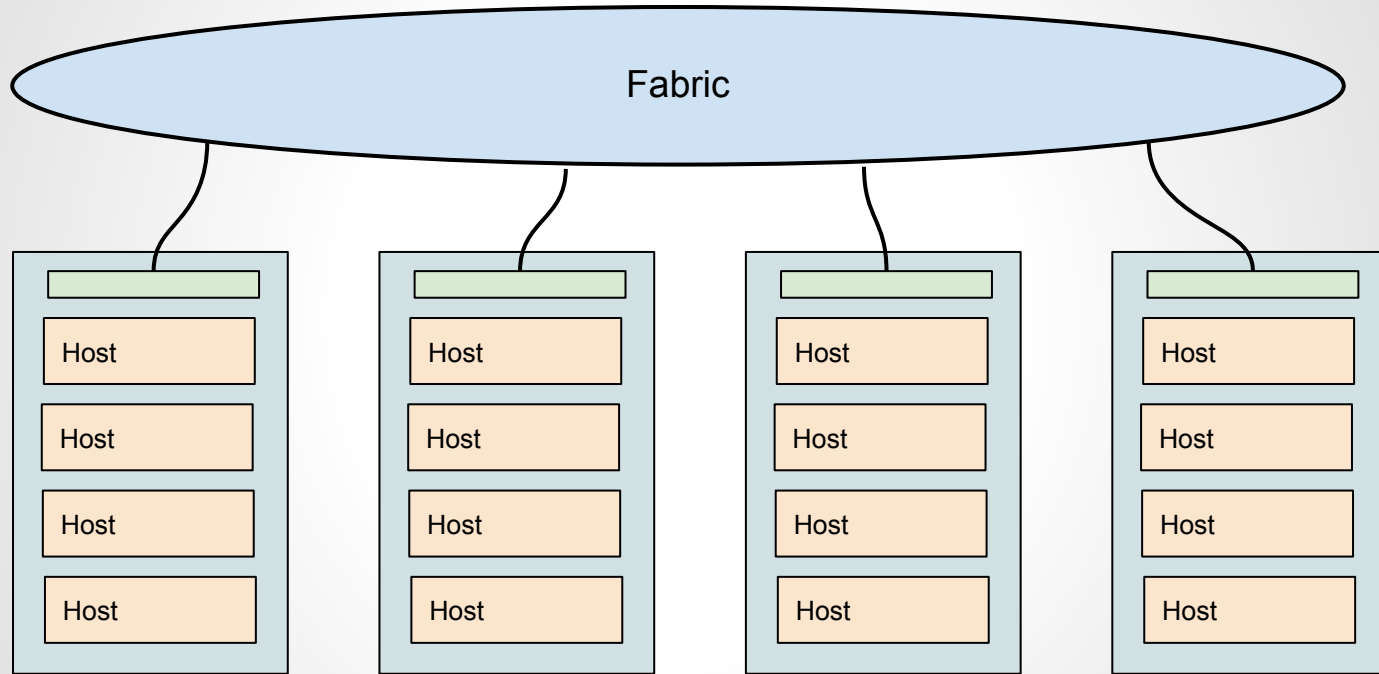
Motivations

- Task Virtualization
- Connectivity for VMs (external to VN)

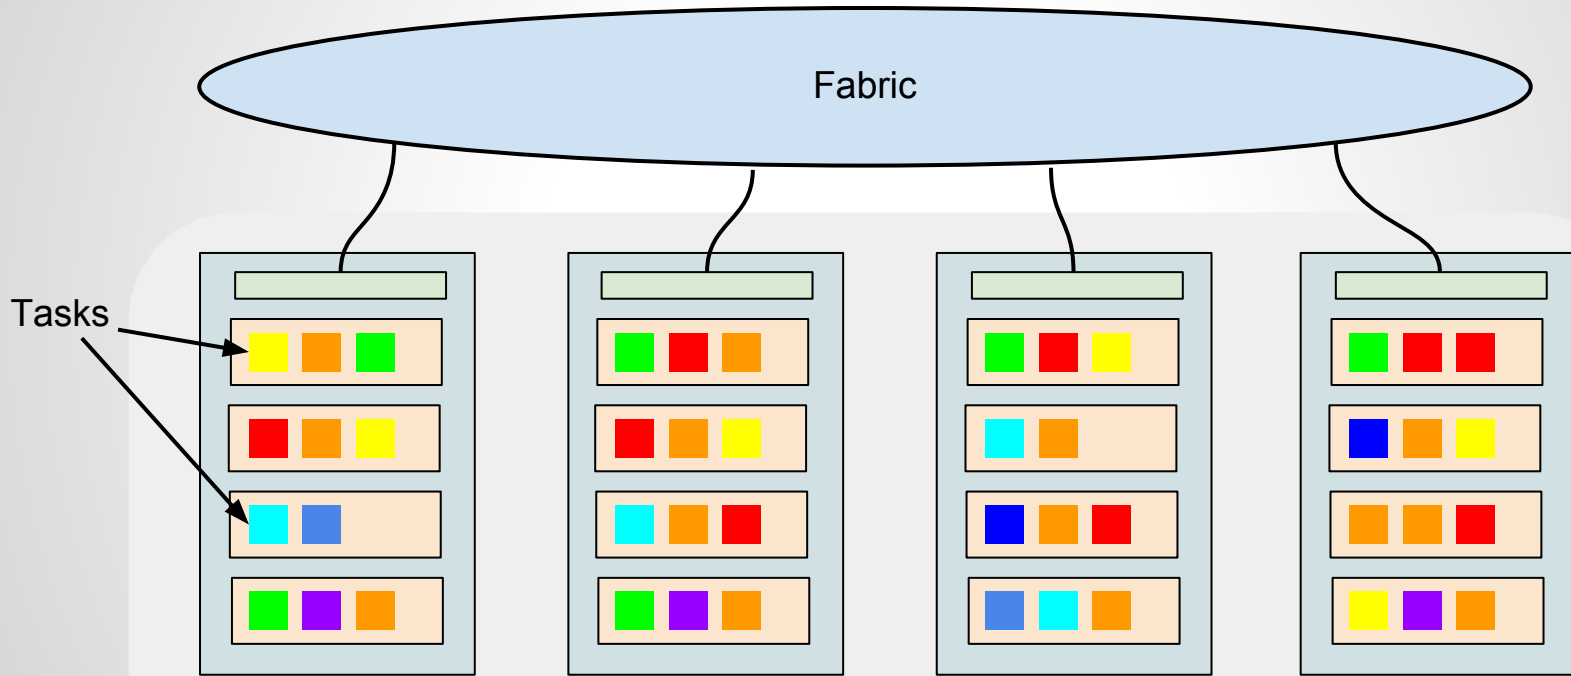
Task virtualization

Capability that every task in the data center can be seamlessly live migrated per discretion of a job scheduler.

Data center topology

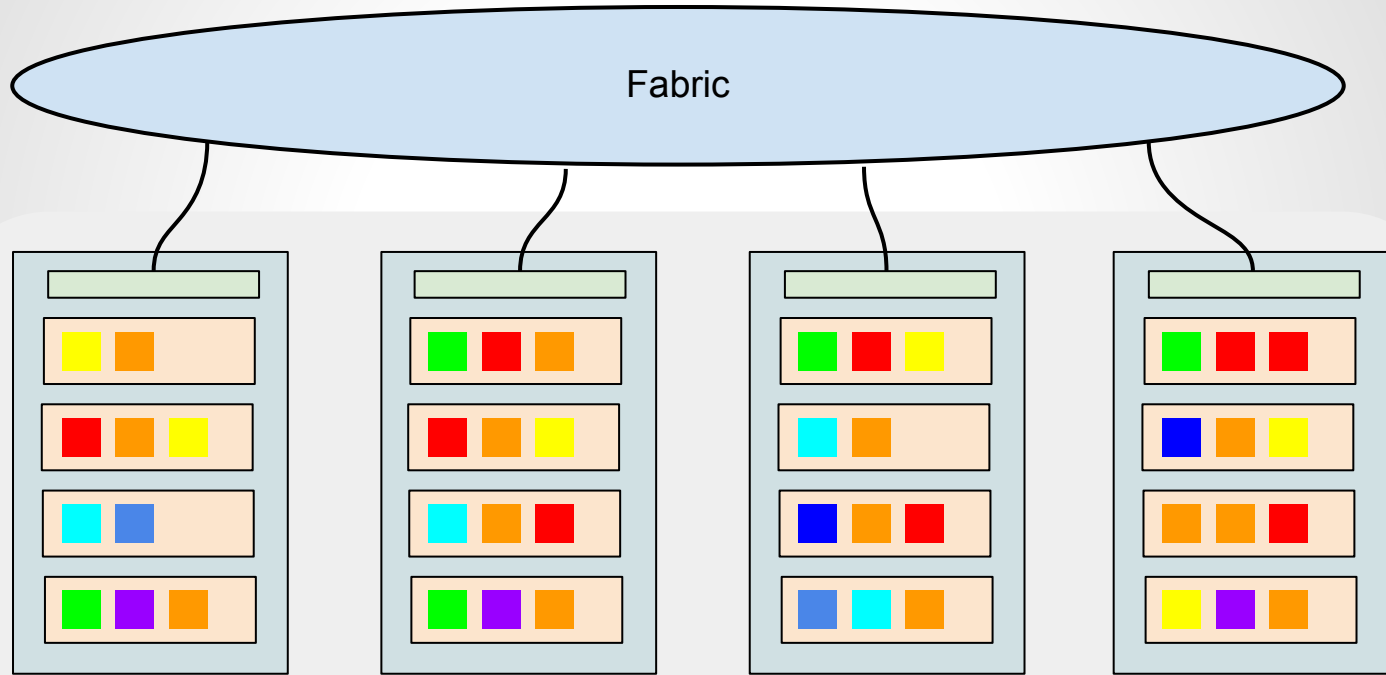


Job scheduler



Job scheduler schedules tasks for each job base on resource requirements

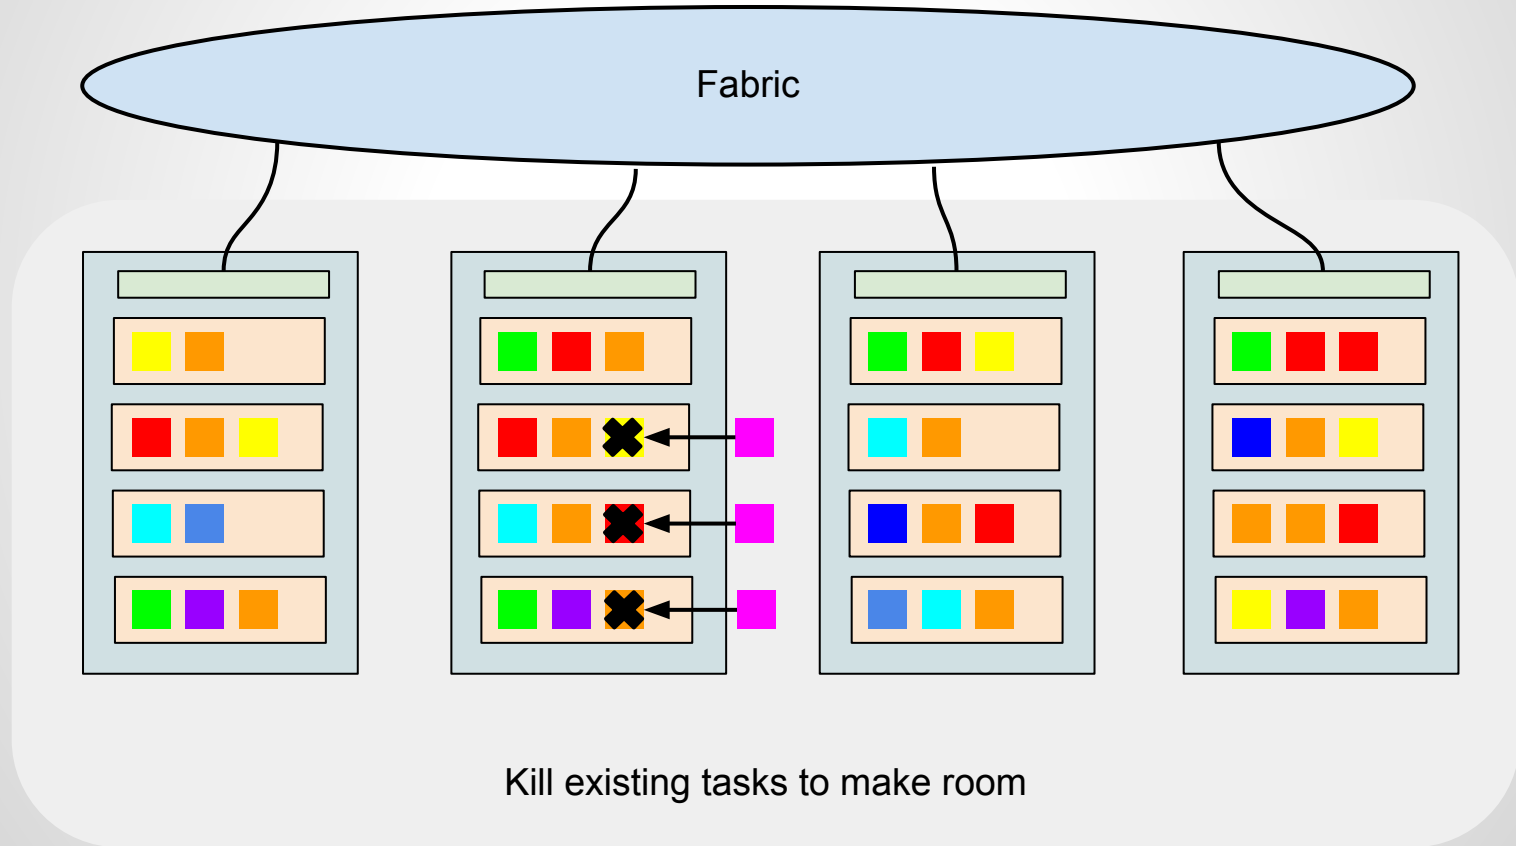
Scheduling dilemma



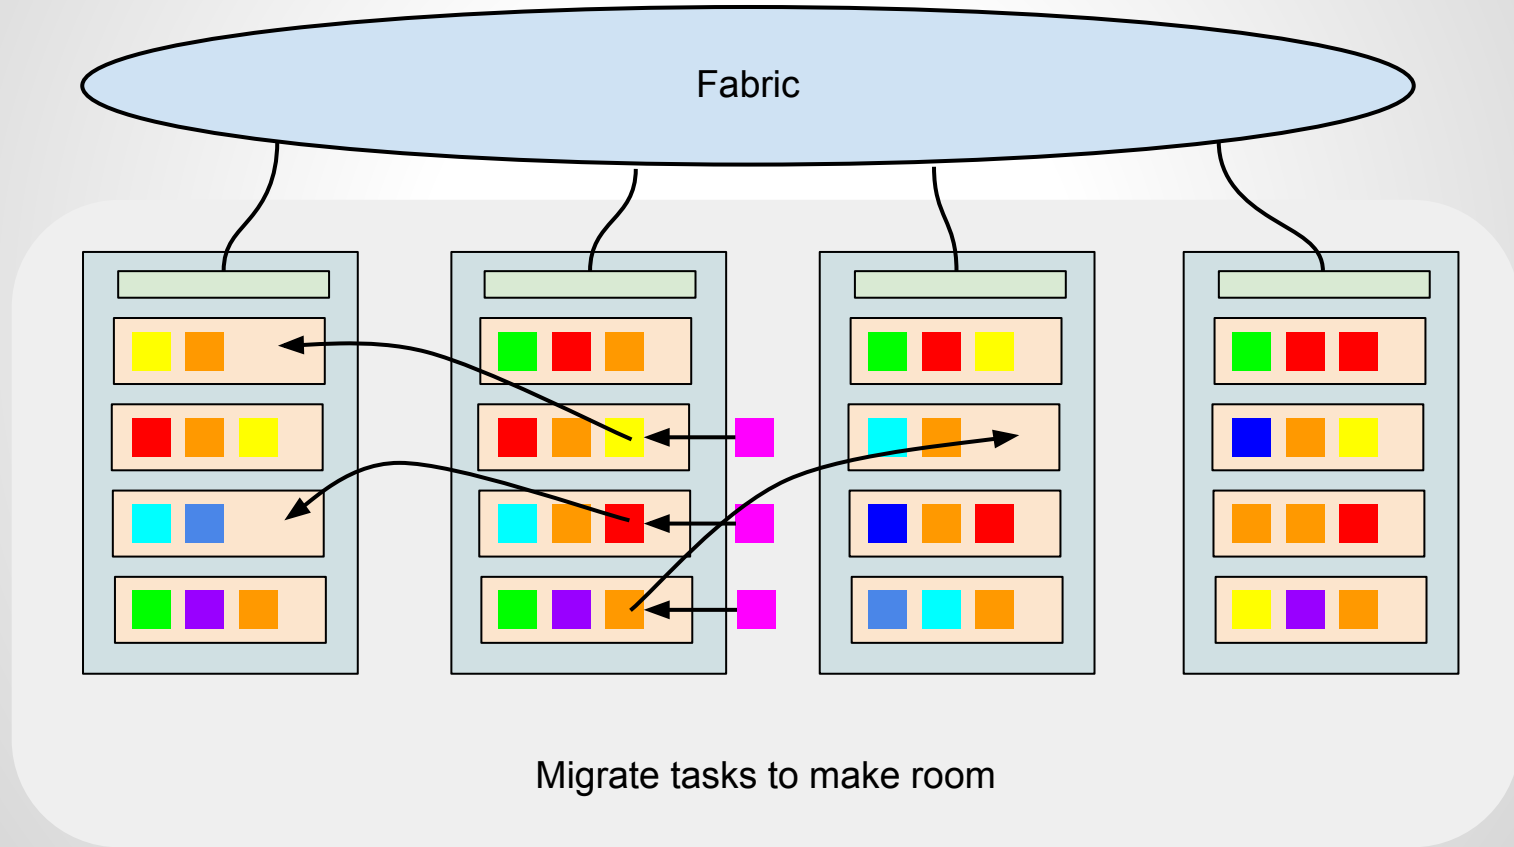
Job scheduler: new, high priority
job with resource constraints



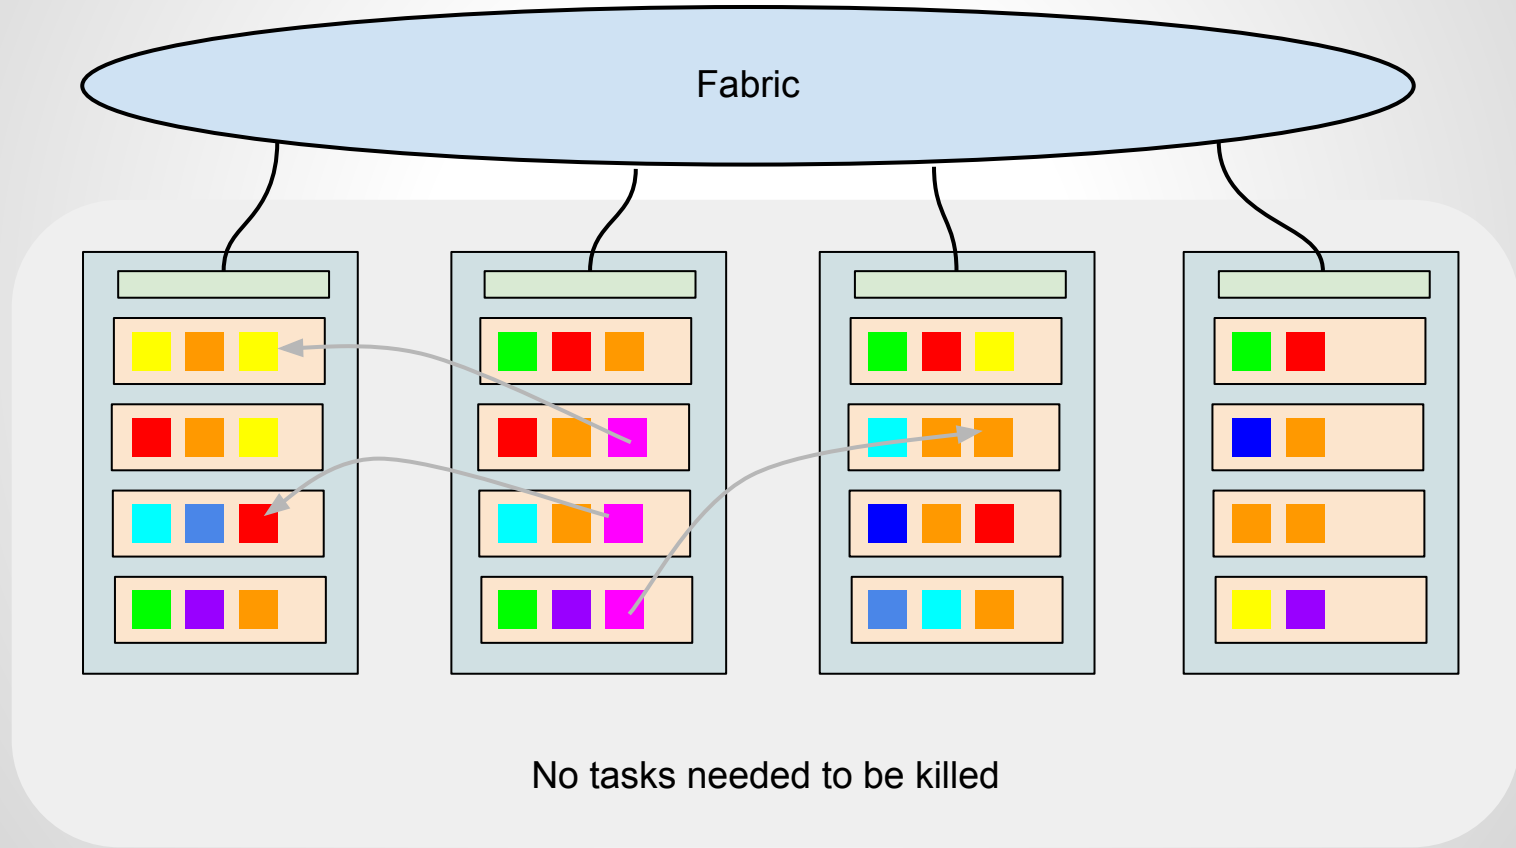
Unpleasant solution today



Task migration solution



After migration




Requirements/assumptions

- Be **transparent** to apps, users, & network
- **Zero** performance impact when not migrating
- Does **not** adversely impact security or control
- Tasks run in containers not VMs
- No overlay networks, no vswitch needed
- ECMP and NIC offloads continue to work
- Most tasks will probably never be migrated

ILA Solution

- Split IPv6 address into identifier (who) and locator (where) ala ILNP
- Each task gets its own unique identifier
- Mapping identifiers to locators
- If task migrates between hosts, its locator changes but its identifier does not
- When not migrating, data path is essentially same as before

Address split



- **Locator**
 - 64 bits identifier of physical hosts
 - Routable
 - Not used as connection endpoint
- **Identifier**
 - 64 bit logical endpoint address of virtual node
 - Not routable
 - Used as connection endpoint
 - Typed to allow different mode

Network virtualization use cases

Locator	Type	VNID	Vaddr
---------	------	------	-------

- Embed VNID in ILA address
 - Potentially eliminate encapsulation for NVO3
 - No place to put security to authenticate VNID, so intra VN use might be limited
- Allows VM to common DC service, or Internet w/o stateful NAT or encapsulation
- Allow two VMs to communicate under policy w/o NAT

Details

- Need to map identifiers to locators
 - Same problem of mapping Vaddr to Paddr in NV
 - Use NVO3 control plane to distribute mappings
- User visible representation of ILA address (used to connect, in DNS, etc.)
- Generate identifier for each task, duplicate address detection

Status

- I-D posted
- Plan to integrate into Linux stack
- Linux implementation is (barely) underway