


LISA 2007 - HTGR



Ganeti

an open source high-availability cluster based on Xen

Guido Trotter
Google Ganeti Team

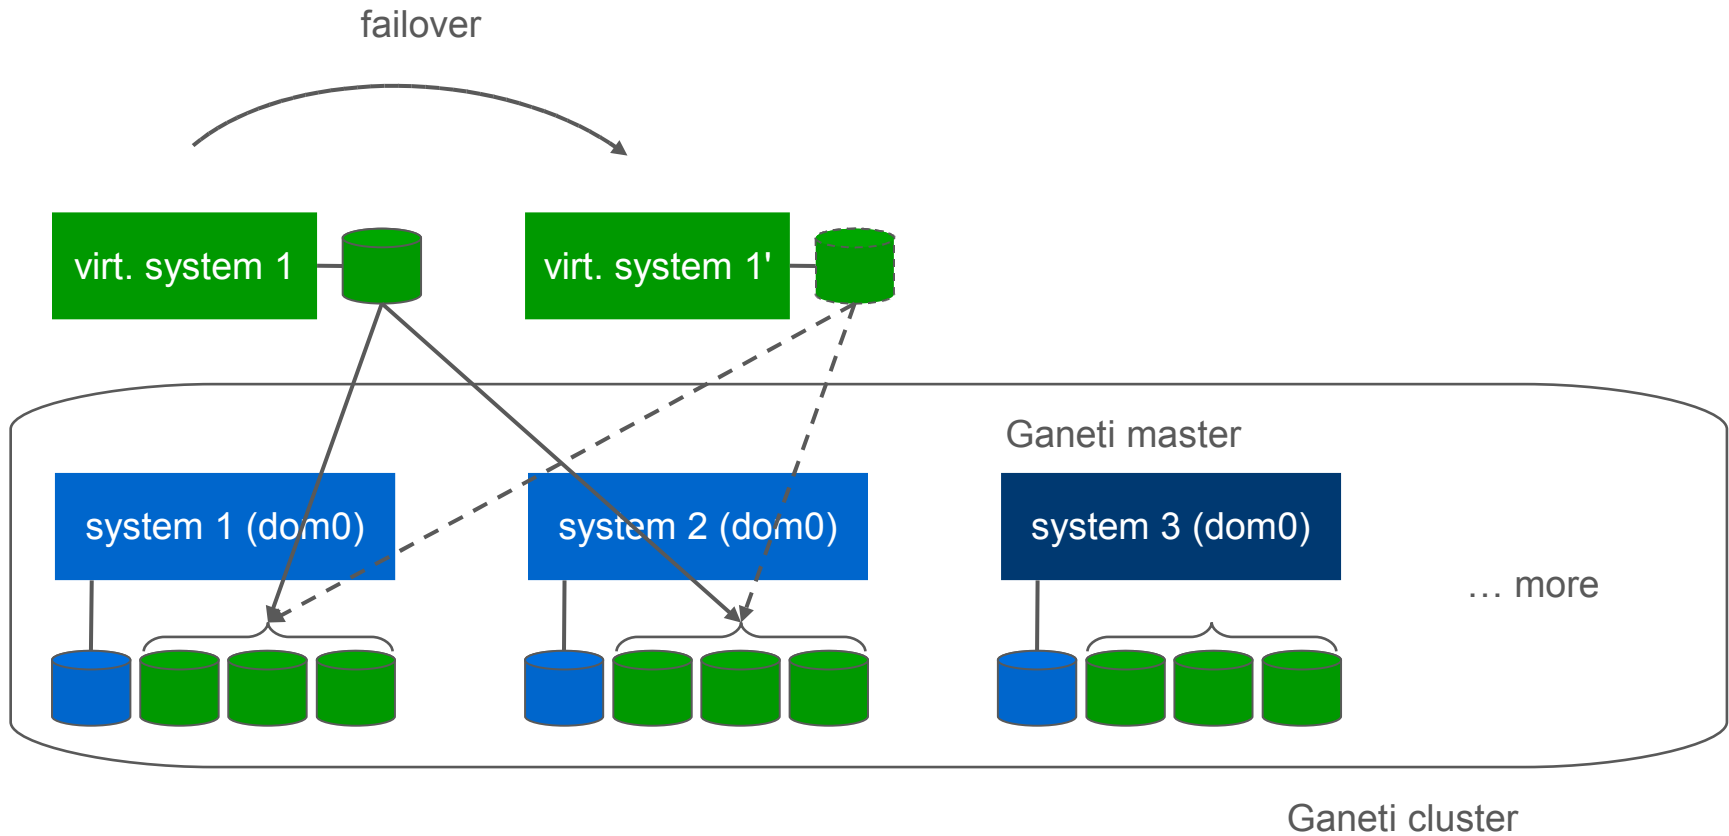
- Design goals and principles
- Ganeti overview and administration
- Ganeti failover details
- Ganeti usage in Google
- Ganeti roadmap

Ganeti is a software to manage clusters of virtual servers

- Based on Xen (but not strictly dependent on it)
- makes it simple to manage 10s of nodes and 100s of instances
- not dependent on specific hardware (e.g. SAN)
- supports different host systems

Ganeti overview (2/2)

Xen dom0 = node
Xen domU = instance



Cluster Setup:

```
node0# gnt-cluster init mycluster
node0# gnt-node add node1
node0# gnt-node add node2
node0# gnt-node add node3
node0# gnt-cluster command \
> apt-get install ganeti-instance-etch
```

Creation of an instance:

```
node0# gnt-instance add \  
> -n node2:node1 \  
> -t drbd8 \  
> instance0
```

Migration after a node crash:

```
node0# gnt-instance failover --ignore-consistency instance0  
node0# gnt-instance replace-disks -s \  
> --new-secondary=node3 instance0
```

Cluster status:

```
# gnt-instance list
```

Instance	OS	Primary_node	Autostart	Status	Memory
instance1.example.com	etch	node1.example.com	yes	running	128
instance2.example.com	etch	node3.example.com	yes	running	512
instance3.example.com	etch	node3.example.com	yes	running	1024
instance4.example.com	etch	node2.example.com	yes	running	128
instance5.example.com	etch	node4.example.com	yes	running	512

```
# gnt-node list
```

Node	DTotal	DFree	MTotal	MNode	MFree	Pinst	Sinst
node1.example.com	858240	442752	4095	511	3456	1	2
node2.example.com	572160	567296	4095	511	3456	1	2
node3.example.com	858240	858240	4095	511	2048	2	1
node4.example.com	356032	356032	4095	511	3072	1	0

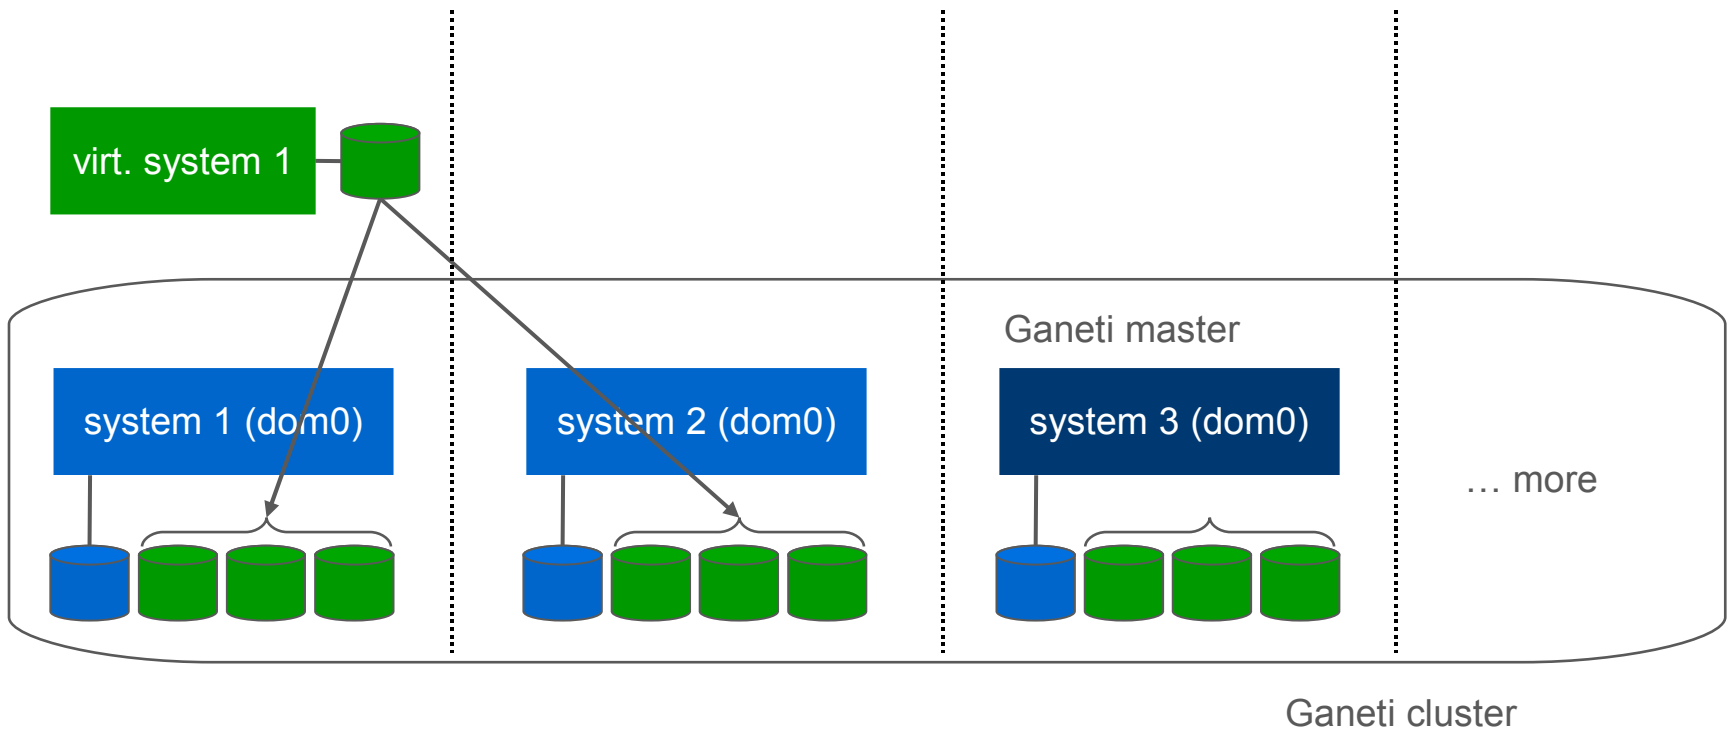
The commands are run on the master node

- `gnt-node`: **add / remove / list** cluster nodes
- `gnt-instance`:
 - **add / remove** instance
 - **failover** instance, change secondary
 - **stop / start** instance, change parameters
- `gnt-os`: **instance OS** definitions
- `gnt-cluster`: **cluster** commands
- `gnt-backup`: **instance export and import**

All commands have man pages and support interactive help.

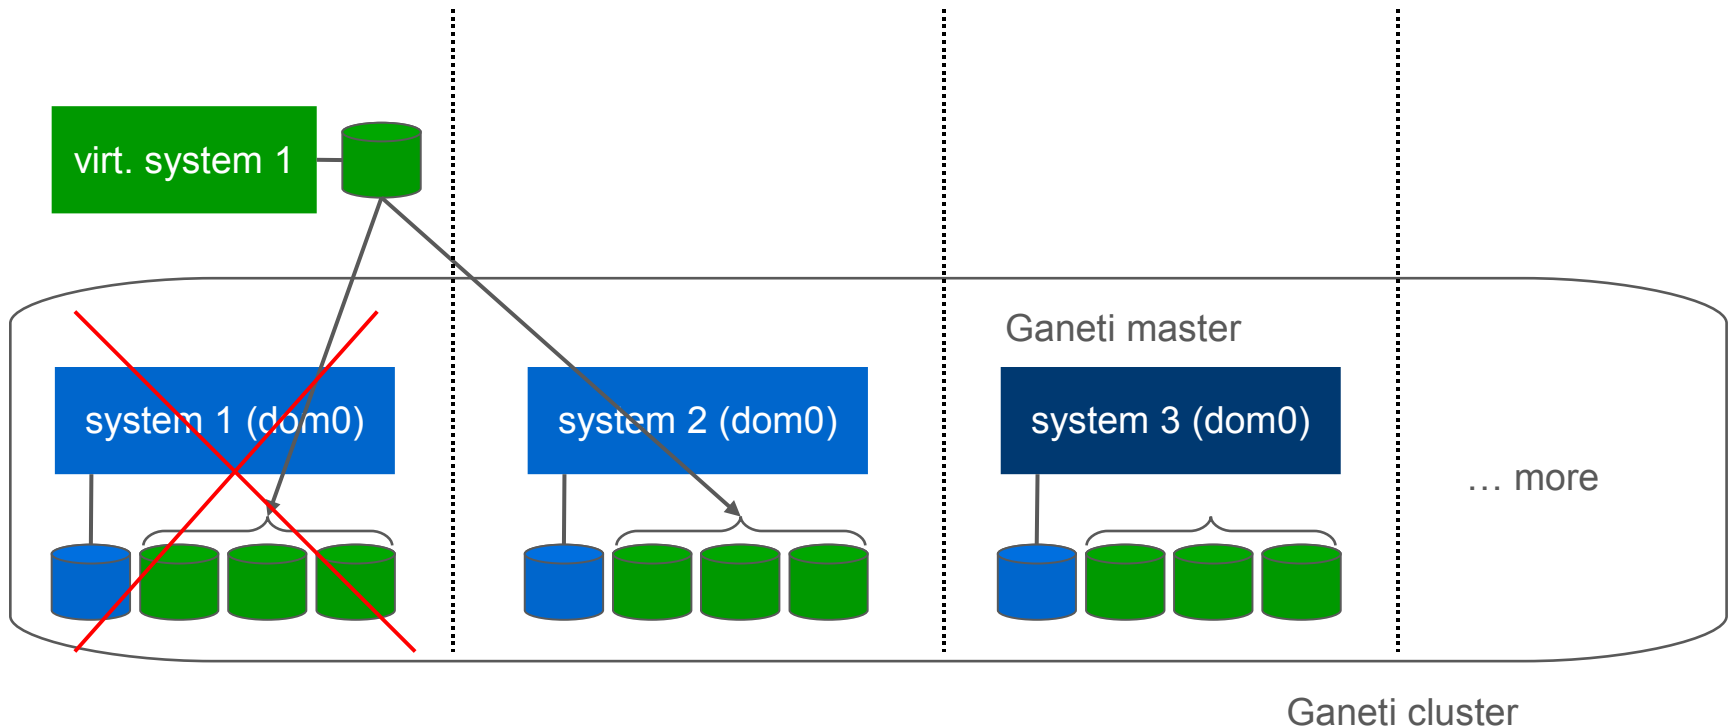
Instance failover (1/4)

Xen dom0 = node
Xen domU = instance

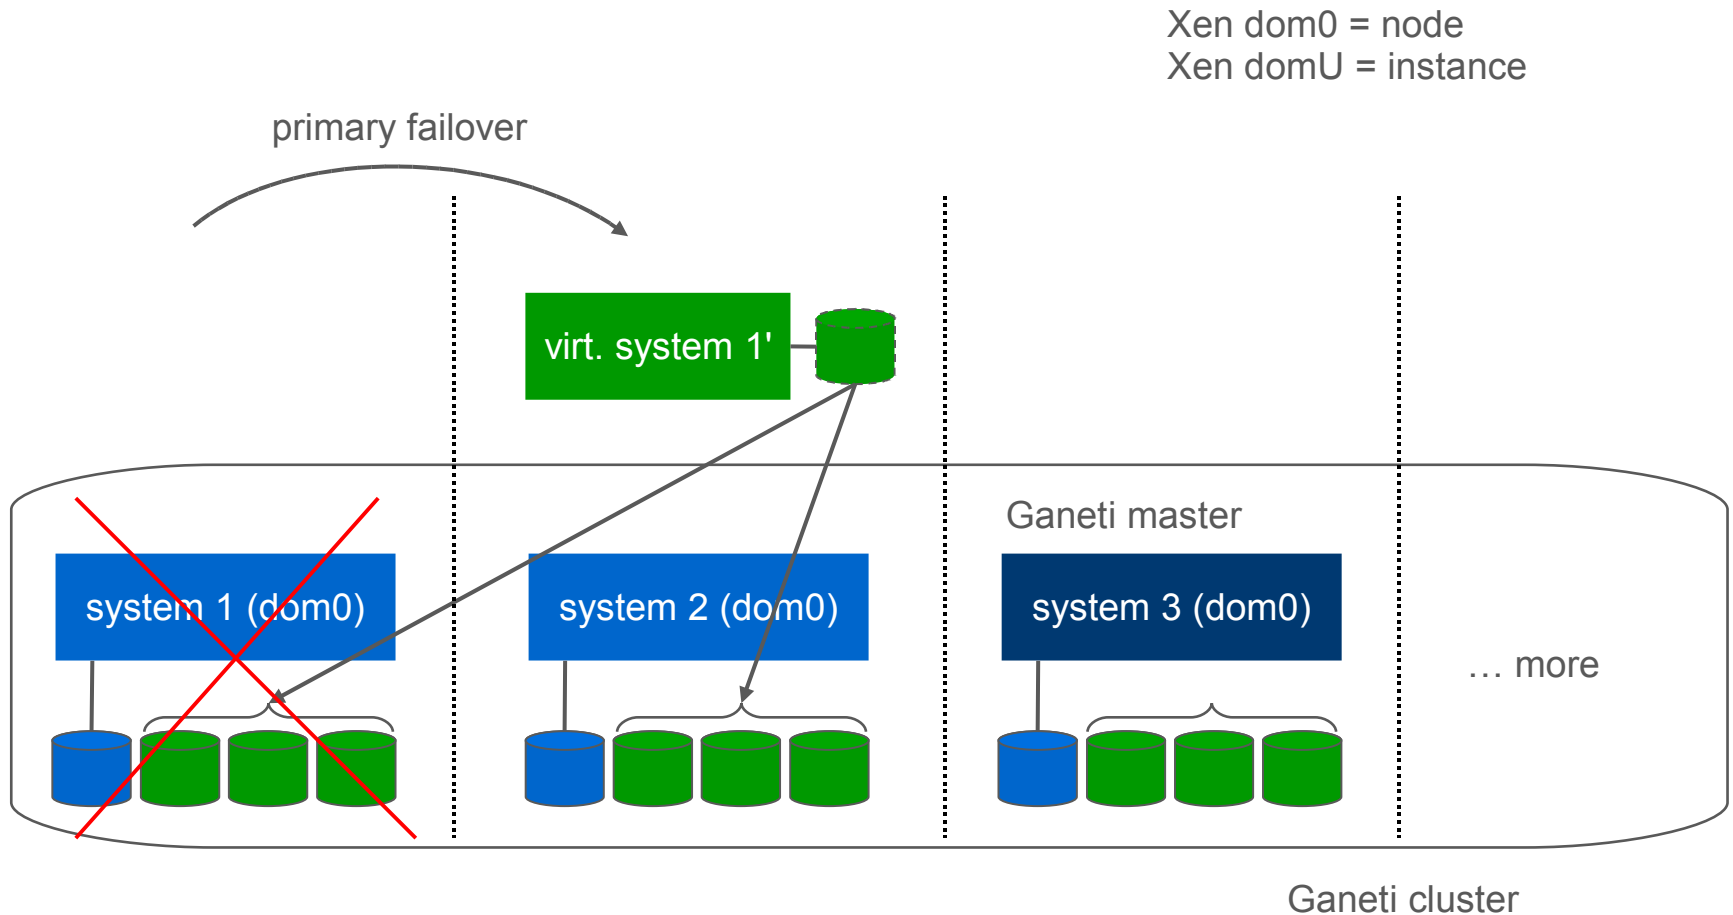


Instance failover (2/4)

Xen dom0 = node
Xen domU = instance

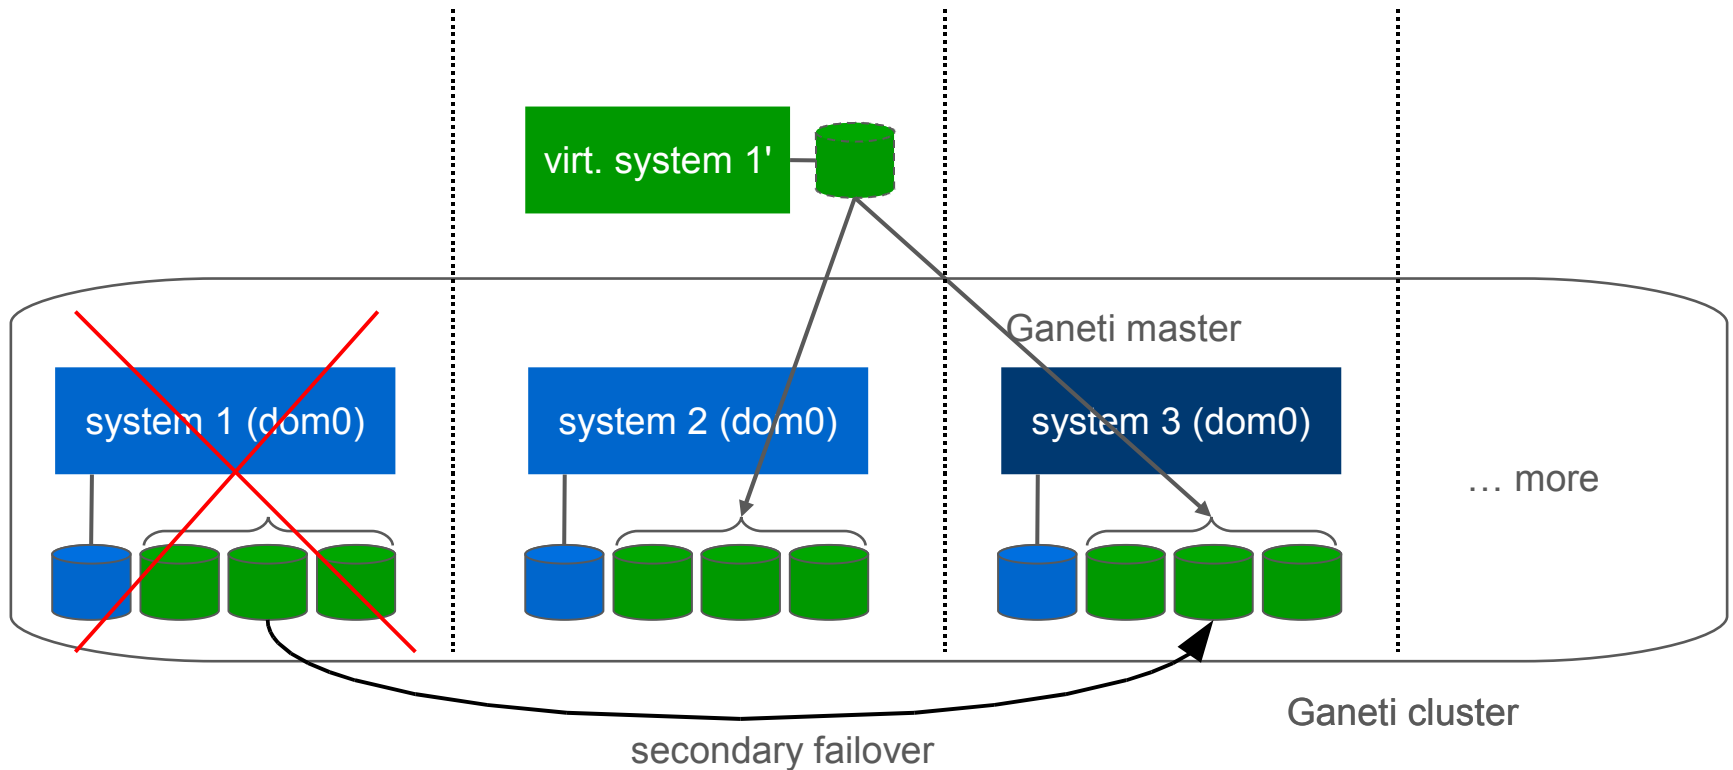


Instance failover (3/4)




Instance failover (4/4)

Xen dom0 = node
Xen domU = instance



Optional advanced features

- Separate replication network
- Multiple bridges/VLAN support
- **Tagging (new)**



- developed at Google
- license: GPLv2
- code location: <http://code.google.com/p/ganeti/>
- August 2007
 - open source and release 1.2b1
- November 2007
 - release 1.2b2
- December 2007
 - release 1.2
- February 2008
 - release 1.2.1
- Later
 - release 1.3 (draft features: remote API, live migration, job queuing, ...)

Q&A

