

minimega

February 2016



Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND NO. 2011-XXXX

UNCLASSIFIED - DRAFT

What To Expect

- Overview and motivation
- minimega toolset
- Use cases
- Current/planned work, goals

The Take

minimega is a research platform that integrates a number of technologies and domains

- Research in Emulytics™
 - Virtualization
 - VM scale/density/scheduling
 - Virtual networking
 - Analytics
- Research enabled by Emulytics™
 - “History has shown that the best way to spawn a ton of research is to provide a tool” – Dean Tullsen
 - MODSIM
 - Analytics
 - RE
 - ...

Emulytics™

- Sandia developed Emulytics™ for cyber security research, testing, and training
- Emulytics™ = emulation + analytics
- Run virtual networks of hundreds of thousands of virtual machines (possibly millions)
- Set up realistic environments that mimic real-world corporate and internet systems
- Integrate Windows, Linux, and Android VMs into a single virtual environment
- minimega is one of several Sandia Emulytics™ tools

A Brief History

- Dissatisfaction with existing tools
 - OpenStack and friends are slow, difficult to use/modify
 - Container implementations aren't well designed for sandbox environments
 - We don't like big, complicated configurations
 - No integrating product available (except OpenStack, see above)
- 2012 – minimega as a simple KVM-based VM management tool
- 2013 – Growth! SDN, command and control, experiment introspection, ...
- 2014 – Android support
- 2015 – Linux container support

Open Source

Full minimega distribution is available publicly (minimega.org)

- Tools, examples, documentation, discussion forum
- Growing community
- Development, issue tracking is done publicly

USE CASES

Test and Evaluation

- Provide a sandbox for “what if” scenarios
- Build experiments around software, hardware
 - Including hardware in the loop
- Provide, run, writeup experiments
- Provide environments, deployment on customer hardware
- Example: DHS Transition to Practice

Training Environments

- Provide detailed environments for red/blue team exercises
- Bridge humans to an experiment
- Support AARs (vnc record/replay)

More Use Cases

- HPC research
- Internet of Things
- Cyber-physical applications
- Honeypot applications

MINIMEGA TOOLSET

Minimega

- minimega the tool vs minimega the platform (nearly 1:1)
- Single binary, deployment is simply copying the binary to all nodes
 - Can actually be simpler!
- Command line, scriptable, programmable interfaces
- Fully distributed (no master/slave semantics)
- No configuration - just launch the tool and start describing your experiment
- No external software stack – no database, no message passing queue, ...

Minimega Script

```
# describe a vm
vm config disk foo.qc2
vm config memory 2048
vm config net my_network

# launch one vm named 'foo'
vm launch kvm foo

# launch 5 unnamed vms
vm launch kvm 5

# grab some info
vm info

# grab some pcap
capture pcap vm foo foo.pcap

# record vnc framebuffer and keyboard/mouse interactions
Vnc record fb foo foo_vnc.fb
Vnc record kb foo foo_vnc.kb


# clean up
vm kill all
vm flush
```

Minimeta At Scale

- Same command set for clustered instances
 - minimeta can union nodes to provide a consistent interface regardless of size and *topology* of the underlying cluster
 - Support for esoteric physical layers
 - Multiple federated clusters, etc.
- Fully distributed
- Cluster orchestration layer
 - **Very fast** file transfer, shell, kernel optimizations (ksm, huge pages)
- Can self-deploy
 - `deploy my_cluster_nodes[2-500] username password`

Minimega Mobile Extensions

- Several radios/sensors simulated in minimega for Android endpoints
 - GSM, wifi, accelerometer, GPS, ...



Protonuke

Zero-configuration traffic generation and protocol servers

- Single binary, zero dependencies
- Windows, Linux, MacOS, FreeBSD, Android
- ARM, x86, x86_64
- HTTP, HTTPS, SMTP, SSH custom server with content generation
- Servers/client behavior and datasets can be changed with optional configuration
- ~30k emails embedded into the binary in 4 languages
- Can saturate 40Gbit links in a single thread
- Can hit billions of transactions / day when distributed
- Very active development

Miniccc

Command and control agent

- Send/receive files (including globs)
- Execute/background commands
- Process control
- Windows, Linux, MacOS, FreeBSD
- ARM, x86, x86_64
- Can connect over TCP, virtio, or UDS
- Virtio, UDS allow networkless backchannels
- Supports TCP tunneling (equivalent to ssh -L and -R)
 - Even over networkless backchannels!

Vmbetter

Debian-based VM image builder

- Simple, nested VM configurations
- Can build qcow2, kernel/initrd, livecd, containers
- Great for building simple images that get provisioned with miniccc
- Suitable for host images (our cluster nodes PXEboot vmbetter images)

Minidoc

Website, documentation, articles, and presentations (except this one)

- Renders markdown-like articles and slidedecks
- `minimiga.org` is running minidoc, combination of HTML and rendered articles
 - The entire `minimiga.org` website is part of the distribution
- Can be run locally
- Can *interface* `minimiga` for interactive articles/presentations
- Runnable examples in documentation


Other Components

- **Tools**
 - **rfbplay**: VNC record/replay framebuffer playback and transcoding tool
 - **igor**: node reservation tool
 - **powerbot**: PDU control tool
- **Custom libraries**
 - **minicli**: Command line parser, API compiler
 - **meshage/iomeshage**: mesh-based message passing protocol, file transfer layer
 - **miniclient**: Go bindings for interfacing minimega
 - **minitunnel**: TCP over any I/O transport
 - **vnc**: Custom VNC/RFV server and client

TECHNICAL DETAILS

Networking - Distributed Minimega


- Depends on a capable physical network
 - 802.1Q VLAN tags (layer 2)
 - GRE/VXLAN (authentic layer 2 with ~30% performance hit)




Meshage

- Mesh-based message passing for all minimega communication
- Auto-discovery
- Node-to-node and broadcast messages
- Mesh structure is resilient to node failure
- Distributing large files over meshage is fast

A Typical Minimega Mesh



A Minimega Meshage Nightmare



Minimega VMs

- Either KVM/QEMU processes or containers
- Configured with `vm config` commands
 - Launch one or more VM from the current config
- Full custom container implementation
 - Very fast!
- Pause, kill, etc. via KVM QMP or container cgroups
- Hotplug USB/CDROM devices
- Write state for migration/backup
- Screenshot running VMs
- VNC interfaces

Minimega API

minimega is designed to be programmed

- Goal is to provide a simple interface to write tools on top of minimega
- 3 programming interfaces
 - Built-in CLI
 - UDS speaking JSON
 - Meshage
- External bindings (currently python, go)

Web Front End

- Cluster, VM status/info
- Enables accessing any VM in the experiments, across the cluster
 - VNC gateway
- Rich set of graph visualizations