

# PowerNap

## Dynamic Power Management

**Dustin Kirkland**

Canonical

Manager, Systems Integration

Ubuntu Core Developer

kirkland@canonical.com



# What is PowerNap?

- Like a **screen saver**, but for **servers**
- Detects **inactivity**
  - rather than disabling a display, puts underutilized servers into **lower power states**
- Monitors for new **activity**
  - raises servers back to **full power** as necessary
- Ultimately, it's **dynamic power policy management**
- Initially integrated into the **Ubuntu Enterprise Cloud**
- Now, it's a general project/project/solution for **Servers** (and even Desktops)

# Some PowerNap Numbers

System	No PowerNap		PowerNap		
	Busy	Idle	PowerSave / TTR	Suspend / TTR	Hibernate / TTR
Thinkpad x201	35W	16W	<i>13W / 0s</i>	1.9W / ~3s	0W / ~30s
HP 8xCPU 1u Server	430W	300W	<i>280W / 0s</i>	N/A	0W / ~180s

## On a laptop, PowerNap means longer battery

- My normal 4 hour battery lasts over 6 hours with PowerNap

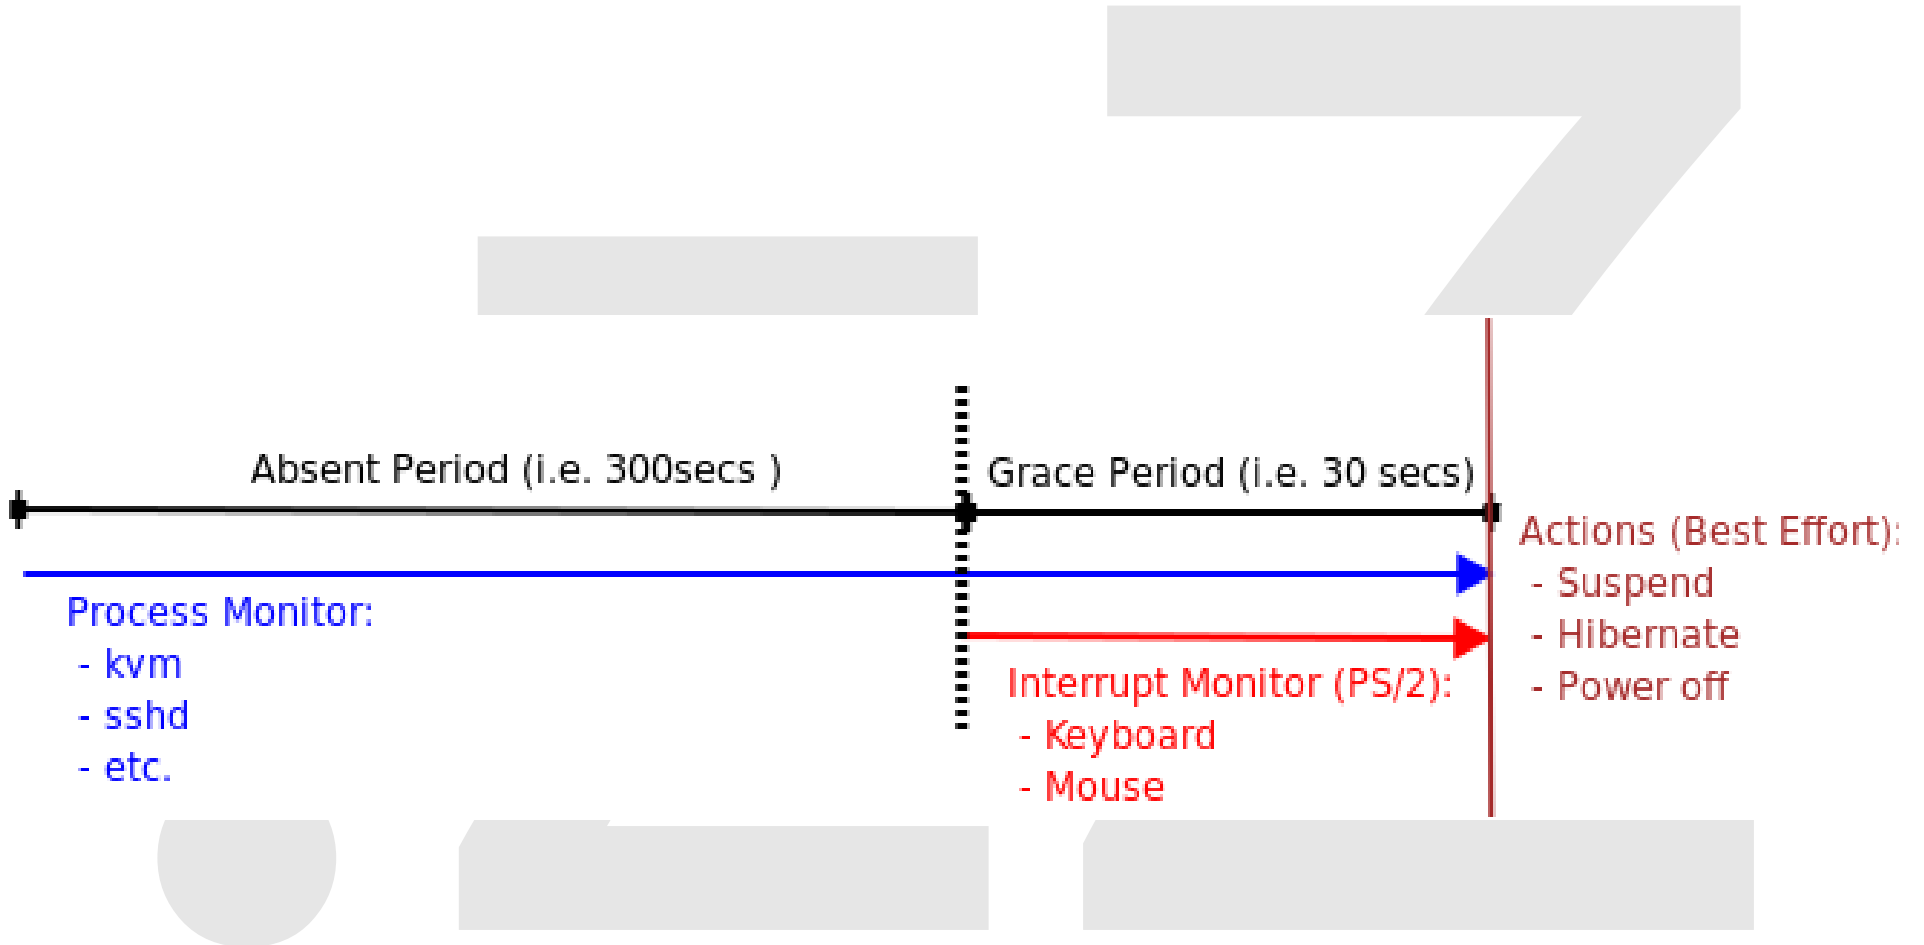
## In a data center, PowerNap means lower energy bills

- At 20W/hour saved, and a rate \$0.10/KWh, that's \$17.52 saved per year, per machine
- Not impressed? What about 10,000 machines x \$17.52 = \$175,200

# The Original PowerNap1 Approach

- **MONITORED PROCESSES**
  - Watch the system process table looking for absent processes
- **INTERVAL SECONDS** (e.g. 1 sec)
  - Interval for which to check for the MONITORED PROCESS
- **ABSENT PERIOD** (e.g. 300 secs)
  - Time for which the process has not been seen
- **GRACE PERIOD** (e.g. 30 secs)
  - Time before performing and ACTION
- **ACTION METHOD**
  - Custom script, Suspend, Hibernate, or Power-off

# PowerNap1 Monitor/Action Timeline



# Motivation: Cloud Integration

- **Eucalyptus**

- SCHEDPOLICY=[ROUNDROBIN,GREEDY,POWERSAVE]  
as a configuration option
- INACTIVITY was tracked by Eucalyptus
- **powernap-now** when a node is running no cloud instances
- **powerwake** nodes when requests exceed capacity of online nodes

# The PowerNap2 Approach

- Make PowerNap **generally applicable** to any Linux data center or server workloads
- Andres Rodriguez's graduate project at FIU
  - Support widely available ways to save power, **without bringing the server entirely offline**
  - Monitor **many different types of activity**
  - In a **highly configurable** manner
  - Fix the **ABSENT/GRACE** period ambiguity

# PowerNap2: PowerSave Action

- Problem
  - Few servers actually support S3/Suspend-to-RAM
  - Hibernate/Poweroff takes a long time to sleep/wake
  - Server is **essentially offline** while in these modes
- Solution
  - **Add a PowerSave state**, to save power while still running
  - **Resume** from PowerSave, and cleanly undo actions
- How
  - **Extend and use pm-utils** power save scripts in  
`/etc/pm/power.d/` on servers



# PowerNap2: PowerSave Scripts

- **Original**, from pm-utils:

- disable\_wol
- hal-cd-polling
- sched-powersave
- intel-audio-powersave
- journal-commit
- sata\_alpm
- wireless

- **New**, from PowerNap:

- cpu\_frequency
- cpu\_online
- eth\_speed
- usb\_autosuspend
- lcd\_brightness

# PowerNap2: New Monitors

- Problem
  - **Monitoring the process table** was not enough
- Solution
  - Extend the ability to **determine a busy or idled system**
- How:
  - **Input/Output** devices
  - **Network** activity
  - **Application** activity

# PowerNap2: Configurable Monitors

- **Input/Output Activity**

- InputMonitor
- ConsoleMonitor
- DiskMonitor

- **Application Activity**

- IOMonitor
- LoadMonitor
- ProcessMonitor

- **Network Activity**

- TCPCMonitor
- UDPCMonitor
- WoLMonitor

# PowerNap2: /etc/powernap/config

## [powernap]

ACTION\_METHOD = 0  
ABSENT\_SECONDS = 300  
GRACE\_SECONDS = 60  
INTERVAL\_SECONDS = 1  
WARN = y  
DEBUG = 0  
STAGE2\_ABSENT\_SECONDS = 0  
STAGE2\_ACTION\_METHOD = 4

## [WoLMonitor]

wol7 = 7  
Wol9 = 9

## [ConsoleMonitor]

ptmx = y

## [ProcessMonitor]

mplayer = "mplayer "  
sshd = "sshd: .\*\[priv\]\$" "  
kvm = "kvm "

## [LoadMonitor]

Threshold = 2

## [TCPMonitor]

ssh = 22  
http = 80  
https = 443  
other = 64500-65000

## [UDPMonitor]

udp = 1025

## [IOMonitor]

kvm-io = "kvm"  
mysqld-io = "mysql"

## [InputMonitor]

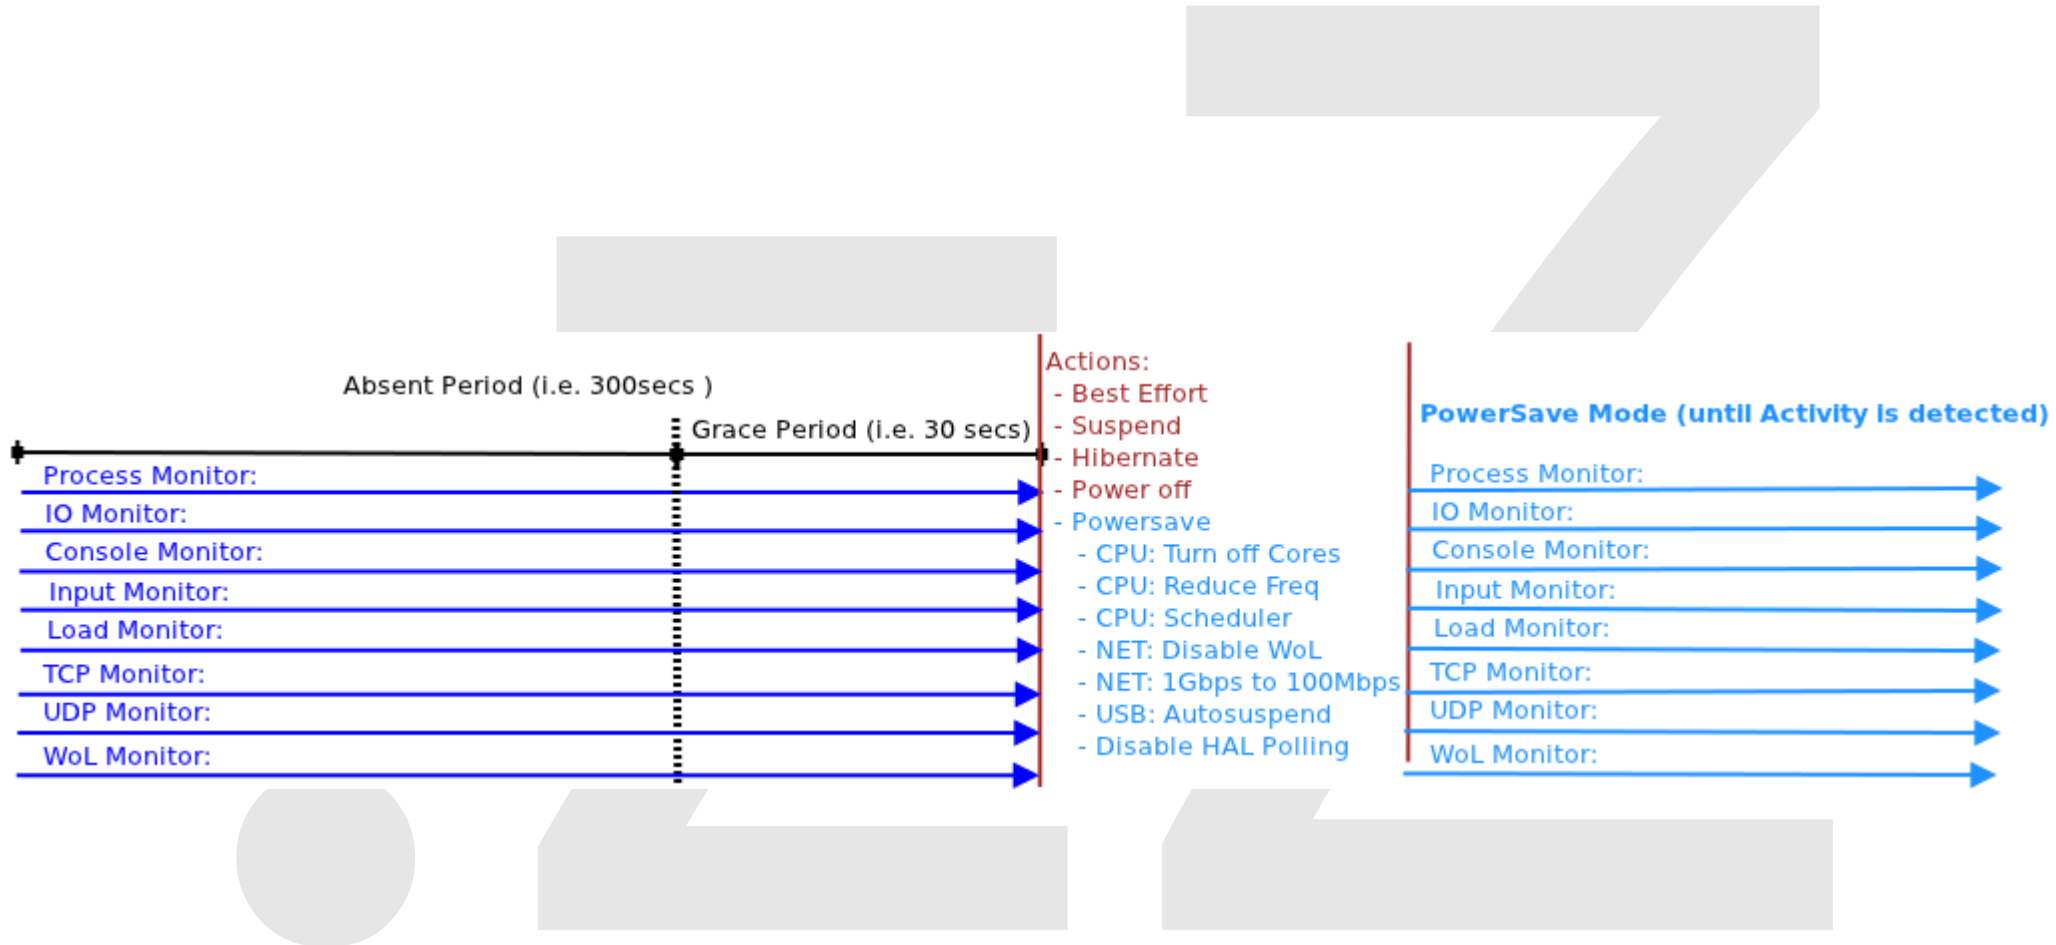
keyboard = y  
mouse = y

## [DiskMonitor]

sda = y

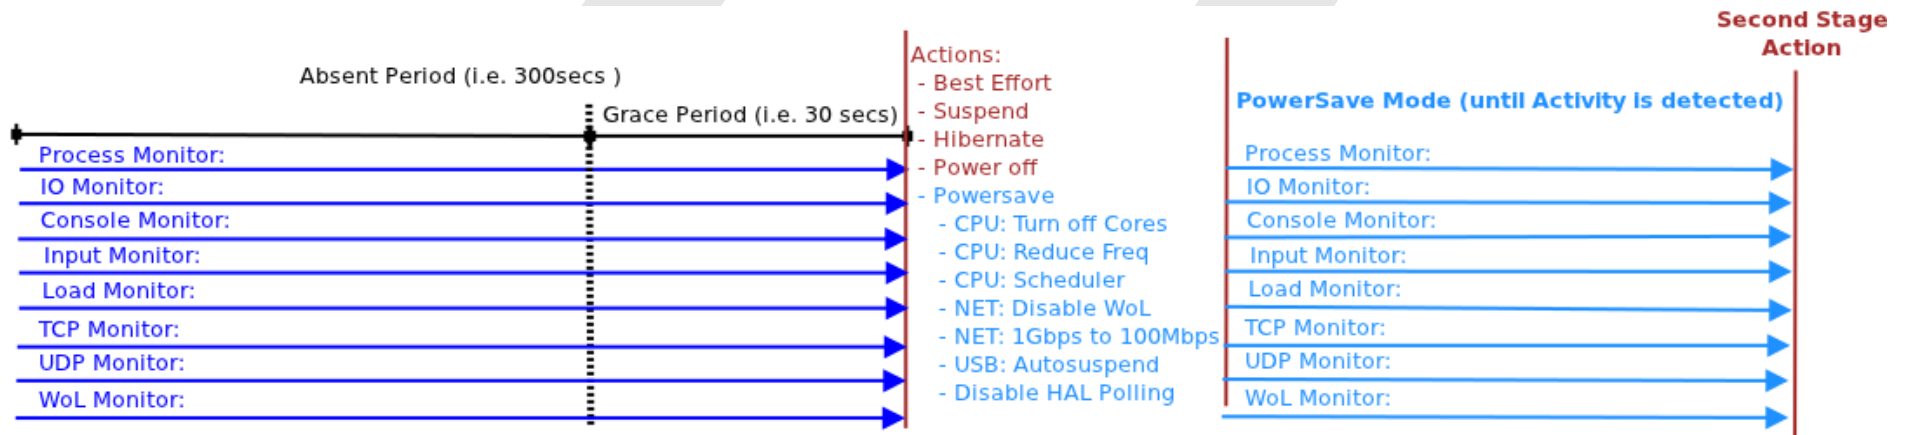


# PowerNap 2.0 Monitor/Action Timeline



# PowerNap2: Second Stage Action

- Optionally take a **second-stage action**, after an extended period in PowerSave state
  - Suspend, Hibernate, or Power-off idled machine



# PowerNap2: Helper Tools

- **powerwake:**
  - Sends WoL packet to IP/MAC address
  - Caches host names, ip addresses, mac addresses
- **powernap-action**
  - Enable/disable action methods for PowerSave
- **powernap-now**
  - Sends a signal to local daemon to execute ACTION
- **powerwake-now**
  - Sends signal to local daemon to recover from ACTION  
(PowerSave)

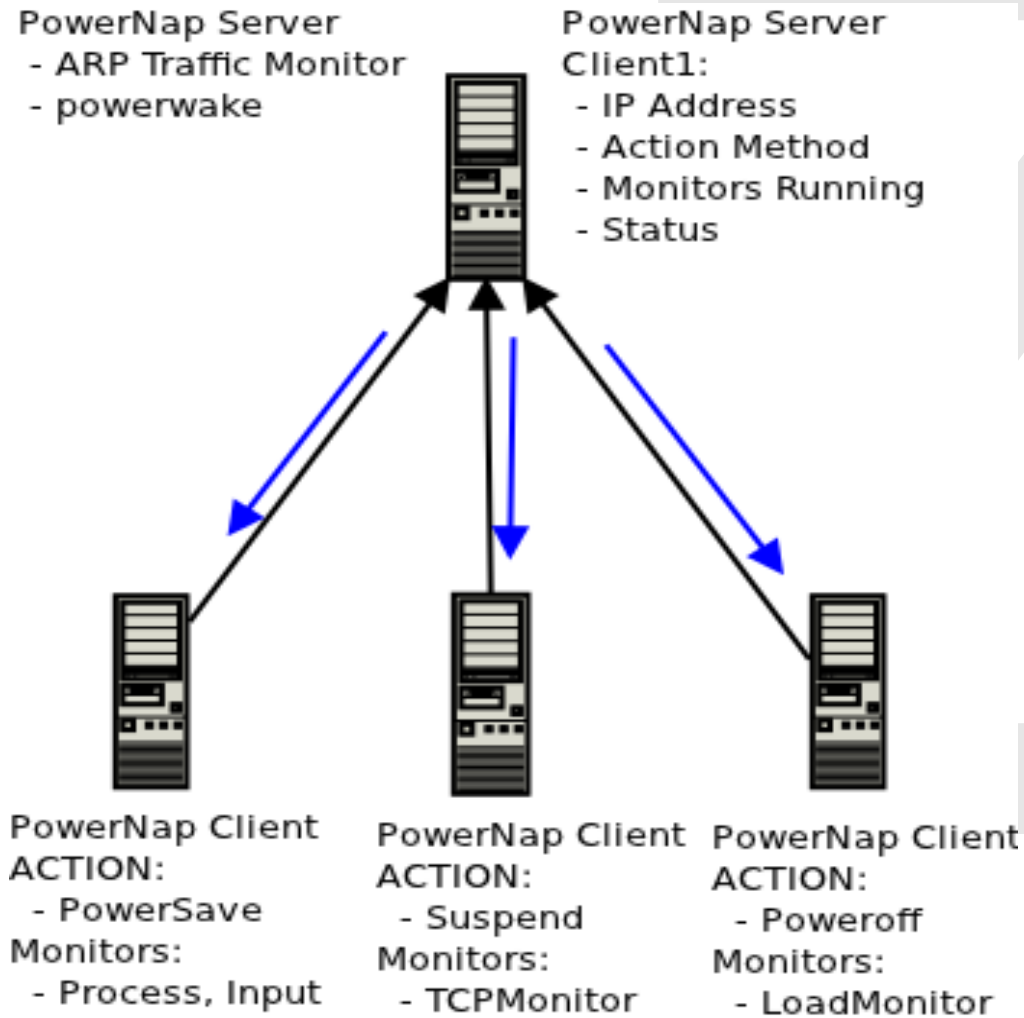
# What's Next?

## PowerNap3: PowerNap Client/Server coming soon!

- Objective:
  - **Manage** machines running PowerNap
- Operations:
  - **Schedule** wake-ups, sleeps
  - **Track** status of machines
  - **Expose** an API
  - **Auto-register** new systems



# Coming Soon: Client/Server Model



# What would we like from Linux Plumbers?

- More enable/disable power savings scripts in **pm-utils**  
`/etc/pm/power.d/*`
  - Ideally, that apply to servers
- More monitors that trigger on other server activities
- Asynchronous monitoring, via `dbus/upstart/systemd`?
- Other distributions?

# Need More Information?

- Website, project, source code, questions, bugs:
  - <http://launchpad.net/powernap>
- **Dustin Kirkland** (original author of PowerNap)
  - [kirkland@canonical.com](mailto:kirkland@canonical.com)
- **Andres Rodriguez** (current maintainer of PowerNap)
  - [andres.rodriguez@canonical.com](mailto:andres.rodriguez@canonical.com)



Questions? Comments?  
Suggestions? Ideas?  
Extensions?

