Questions about ELI

Intro and Background
What is direct assignment of I/O devices?
- What are its costs, what are its benefits?
Why doesn't it alone solve the problem of interrupt overhead in virtualized environments?
Do other techniques help here? (e.g., interrupt coalescing)

Details
What is the IDT?
How does the hardware know its location?
With direct assignment, why can't the host simply let the guest handle I/O interrupts
  for the assigned device?
What is the shadow IDT? What is in it?
What is the NP bit? How does ELI make use of it?
Can a true NP exception occur in the guest? What happens in this case?
How are other device interrupts handled? (e.g., keyboard)
Why is placing the Shadow IDT in memory a challenge?
How does ELI solve this problem?
What is the EOI LAPIC register for?
Why does it cause potential performance problems?
How does ELI solve this problem?
Does the host VMM ever trap accesses to the EOI register?

Threats
What threads are posed by the use of ELI?
(i.e., how could a misbehaving guest cause problems?)
How does ELI address them?
How could better hardware support help?

Evaluation
The setup does not use x2APIC; what problem does this induce for evaluation?
How far off of bare performance is ELI?
How important is using large pages in KVM to back guest memory?
How much does ELI reduce exit rates?
What do Figures 5 and 6 show?