CSI 2008
SECURITY RECONSIDERED

Risks and Challenges of Live Virtual Machine Migration

Presenter:
Jon Oberheide
University of Michigan

CSI COMPUTER SECURITY INSTITUTE

Live VM Migration

- Transfer of a VM from one physical machine to another with little or no service downtime

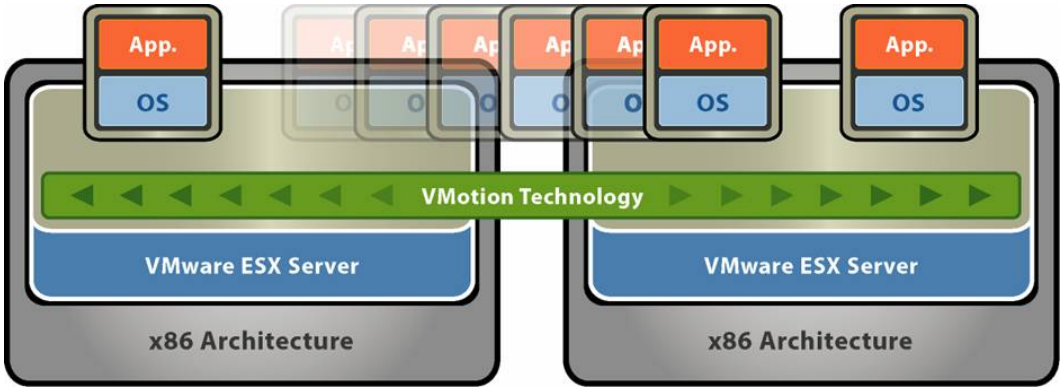

High Availability

Enhanced Mobility

Dynamic Load Balancing

Live Migration Methodology

- Methodology
  - Copy machine state (memory) over network
- Goals
  - Minimize service downtime and migration duration
- Iterative Precopy
  - VMotion (Vmware VI3), XenMotion (Xen Server)

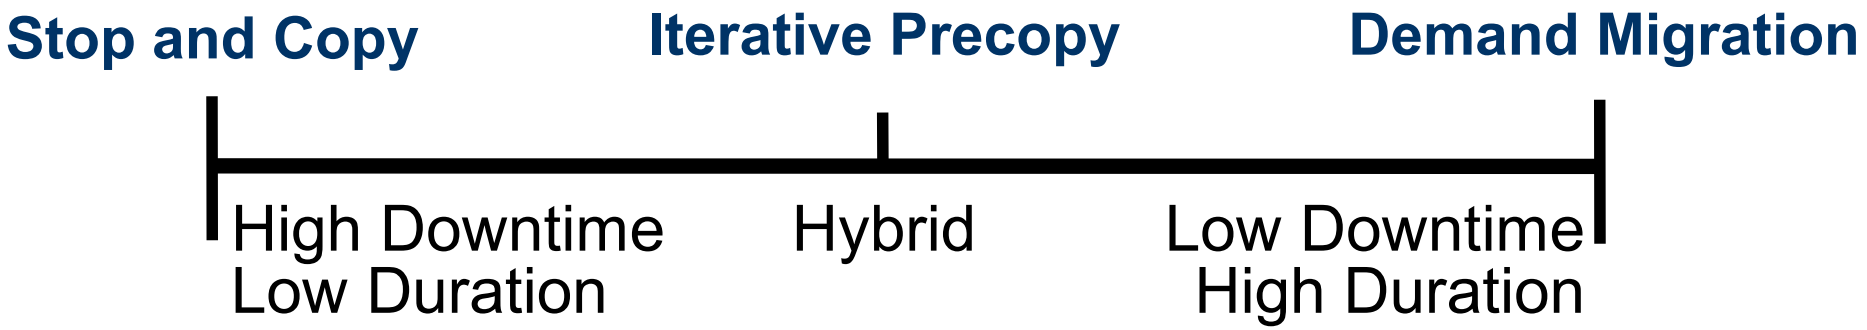

A Trip Down Memory Lane

- Physical machines
  - Machine state protected by MMU/hardware
  - Physical attacks (firewire device DMA)
- Virtual Machines
  - VM state protected by VMM/hypervisor
  - Software attacks (weak VMM isolation)

Can we further weaken isolation boundaries?

Isolation of Machine State

- Virtual machine migration
  - Full VM state exposed to network
  - Trade off of increased risk for functionality and management
  - Authentication, confidentiality, isolation concerns

VM Migration Security

- (In)security of migration protocol
  - Unauthenticated, insecure migration data plane
  - VMotion/XenMotion susceptible to MITM attacks
- Full access granted to VM state
  - OS/kernel memory
  - Application state

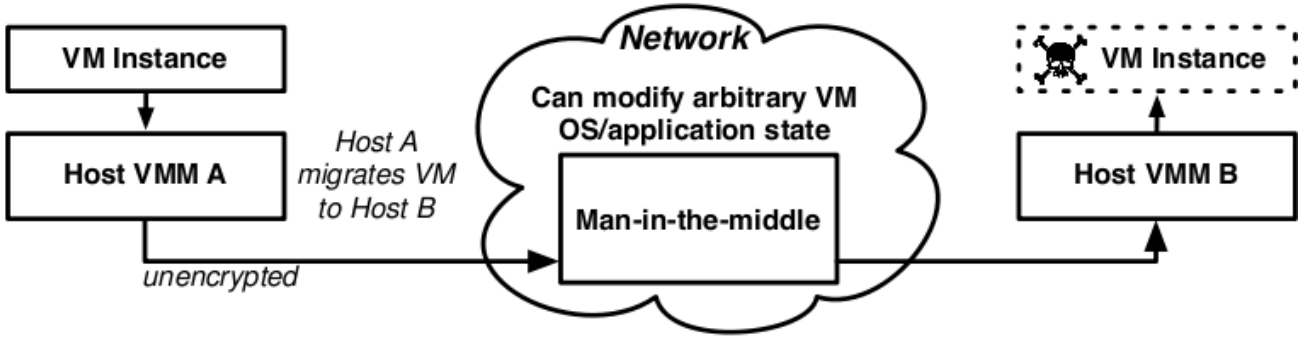

Exploiting VM Migration

- Passive Attacks
  - Snarf sensitive data, passwords, keys in memory
- Active Attacks
  - Manipulate auth. services
    - sshd, /bin/login, etc
  - Manipulate kernel structures
    - slip rootkits into memory

```
        if (key != NULL)
                key_free(key);
        xfree(pkalg);
        xfree(pkblob);
#ifdef HAVE_CYGWIN
        if (check_nt_auth(0, authctxt->pw) == 0)
                authenticated = 0;
#endif
        return authenticated;
}

/* return 1 if user allows given key */
static int
user_key_allowed2(struct passwd *pw, Key *key, char *file)
{
        char line[SSH_MAX_PUBKEY_BYTES];
        int found_key = 0;
        FILE *f;
        u_long linenum = 0;
        struct stat st;
        Key
        ch
        /*
        tem
        deb
        /*
        if
```

```
 805da77:       0f 84 23 fd ff ff       je     805d7a0 <user_key_allowed2+0x80>
 805da7d:       89 3c 24                mov    %edi,(%esp)
 805da80:       e8 37 e5 fe ff          call   804bfbc <fclose@plt>
 805da85:       8d 85 8c df ff ff       lea    0xffffdf8c(%ebp),%eax
 805da8b:       89 44 24 04             mov    %eax,0x4(%esp)
 805da8f:       c7 04 24 15 0e 08 08    movl   $0x8080e15,(%esp)
 805da96:       e8 d5 28 01 00          call   8070370 <logit>
 805da9b:       e8 20 bd 01 00          call   80797c0 <restore_uid>
 805daa0:       81 c4 9c 20 00 00       add    $0x209c,%esp
 805daa6:       31 c0                   xor    %eax,%eax
 805daa8:       5b                      pop    %ebx
 805daa9:       5e                      pop    %esi
 805daaa:       5f                      pop    %edi
 805daab:       5d                      pop    %ebp
 805daac:       c3                      ret
 805daad:       8d 76 00                lea    0x0(%esi),%esi

0805dab0 <user_key_allowed>:
 805dab0:       55                      push   %ebp
 805dab1:       89 e5                   mov    %esp,%ebp
```

VMBR Hoisting

- Virtual Machine-Based Rootkits
  - Slip in extra virt layer (a la SubVirt/Blue Pill/Vitriol)

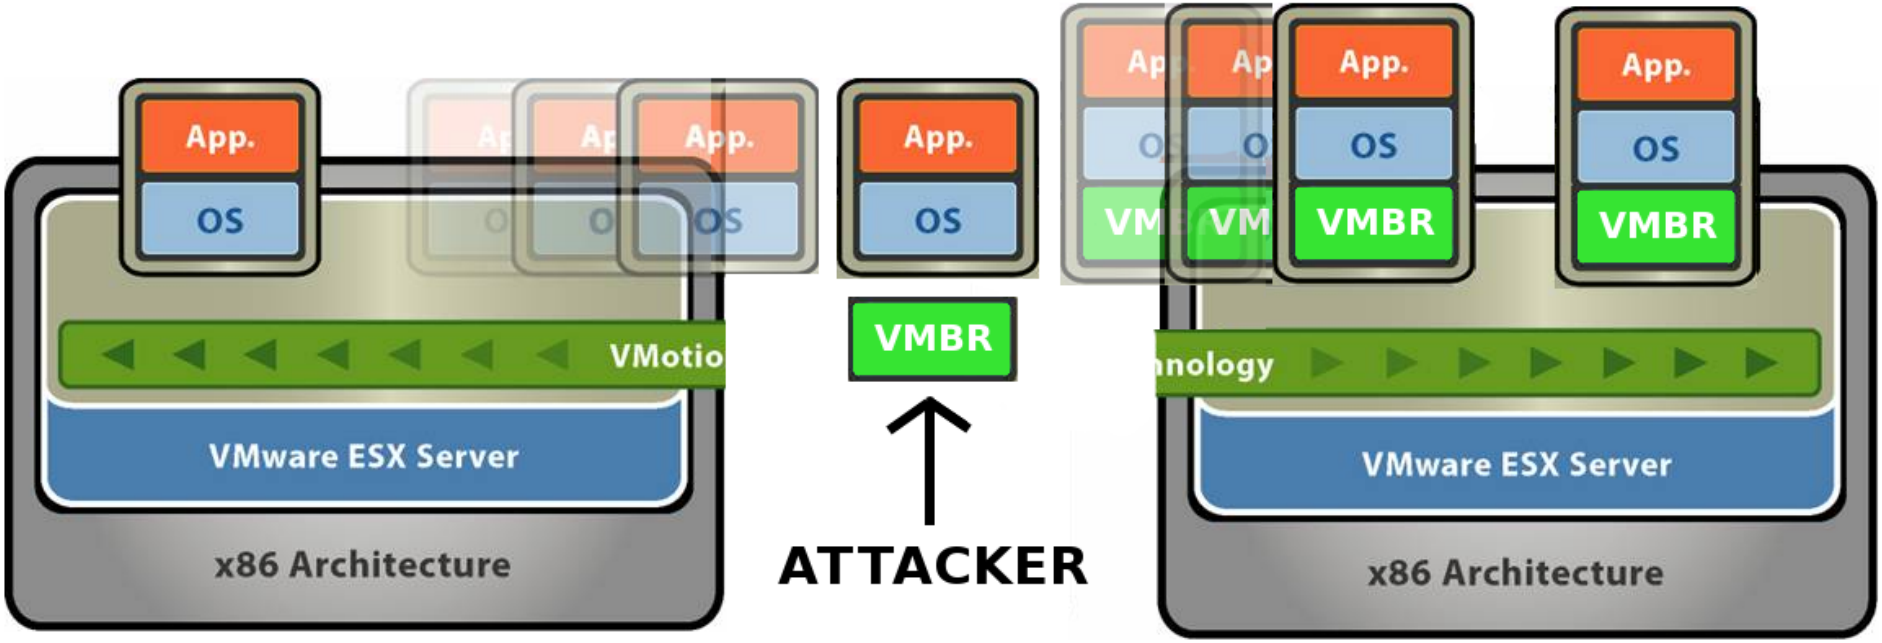

Addressing the Risks

- Encrypt it?
  - Requires authentication to ensure integrity
  - PKI adds deployment and key management complexity
  - Not implemented by vendors
- Isolate it?
  - Separate networks for migration data
  - Physical or virtual (VLAN segmentation)
  - Recommended by VMware best practices guide

Lessons Learned

"Ok, I configured VLANs for isolation, hacked in PKI/TLS support myself, and trained kittens to migrate my VMs.

Am I still at risk? Is this even important any more?"

- Yes and no:
  - Isolation not a *feature* of virtualization, it's a *challenge*
  - Beware of hidden risks in new functionality
  - Best practices and configuration audits are key

Questions?

Q&A

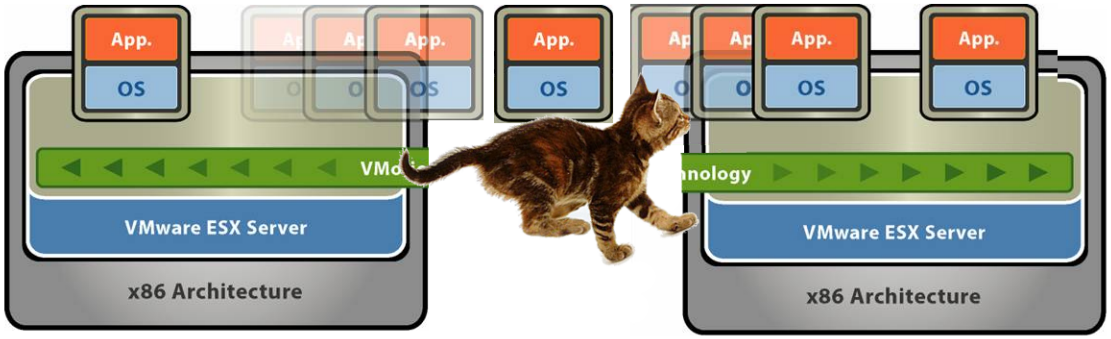

- Contact info:
  - Jon Oberheide <jonojono@umich.edu>
  - PhD candidate, University of Michigan
  - Advisor: Farnam Jahanian
  - Research Group: http://www.eecs.umich.edu/fjgroup/