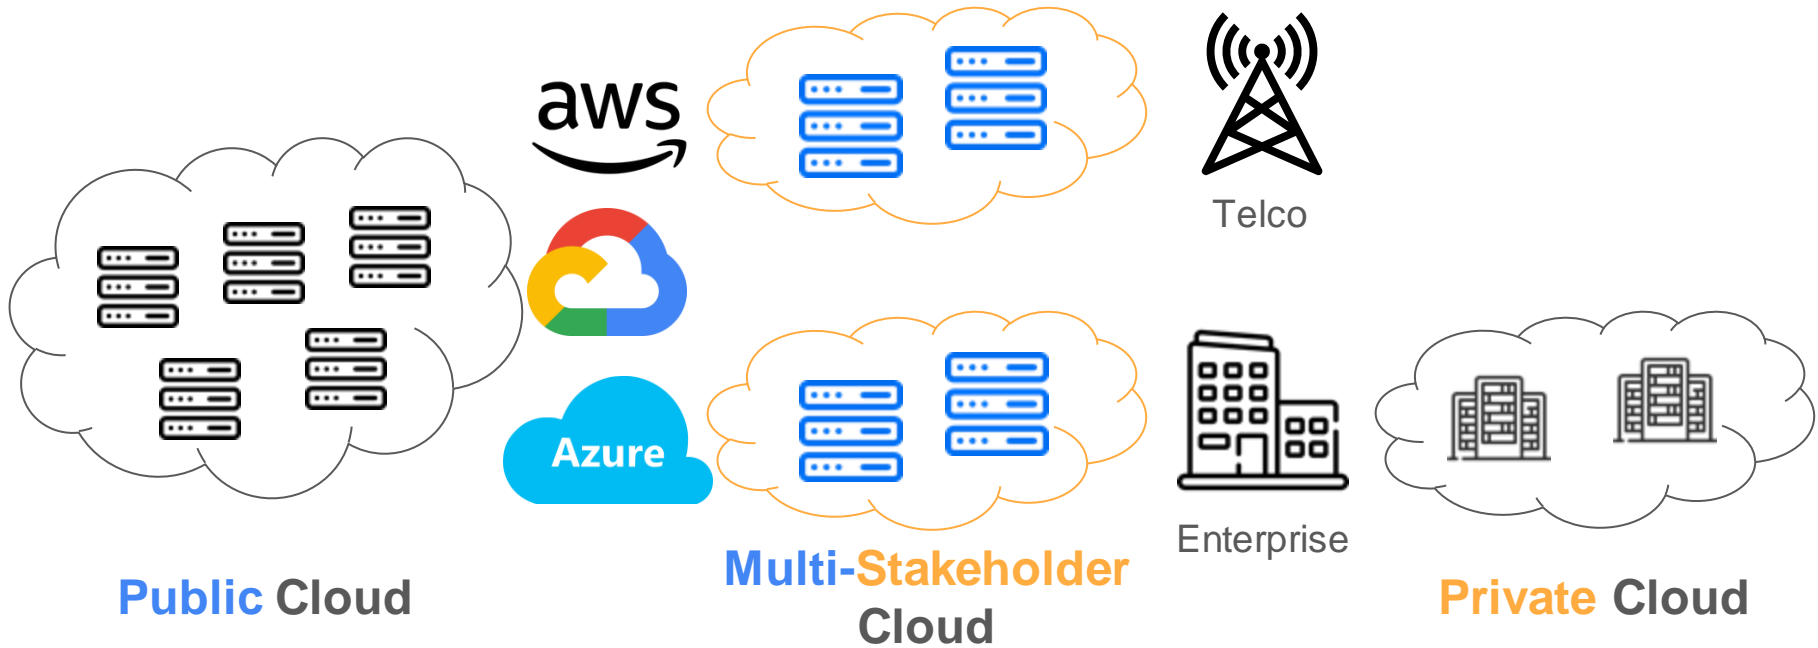


Towards Multi-Stakeholder Clouds

Bohdan Borysei¹, Stefan Saroiu², Eyal de Lara¹


¹University of Toronto, ²Microsoft



Infrastructure Management Tasks

- Power on/off/reboot
- Firmware upgrade
- Provisioning
- Sensors monitoring
- Hardware monitoring
- Events logging
- Networking configuration
- TPM configuration
- Host configuration
- Troubleshooting


Infrastructure Management Challenge due to Multi-Stakeholders



Cloud Provider

Hosting Site Operator









Infrastructure Management Challenge due to Multi-Stakeholders



Cloud Provider

Hosting Site Operator


MCloud Today

	Cloud Provider	Hosting Site Operator
AWS Outpost/Wavelength		
Azure Stack Edge		
Azure Operator Nexus		
Google Anthos		

New Approach Based on Collaboration

*MCloud needs a **new** management primitive that
enables multiple stakeholders
to **collaboratively** manage their infrastructure*

Two-Person Control (P²C)



Two-Person Control (P²C)

- No one party can **independently** perform infrastructure task
- Disperses control, responsibility, accountability **btw** parties
- Used when actions have **significant** consequences
in case of mistakes:
 - Financial, military, pharmaceutical sectors

P²C Authorization Framework Categories

1. Either-party

- *e.g.*, creation/termination of a server management session

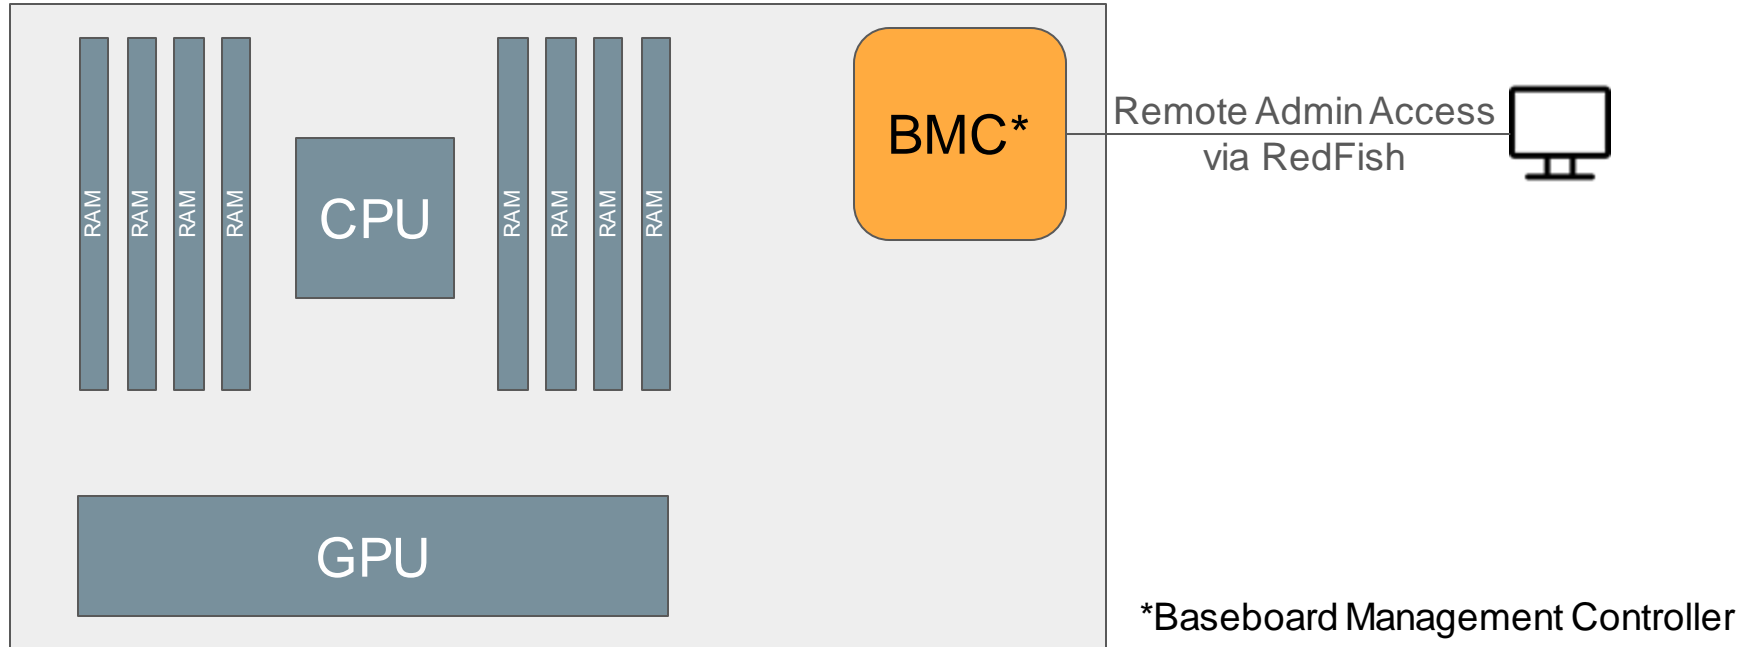
2. Both-parties

- *e.g.*, system factory reset

3. Both-parties-with-policy

- *e.g.*, filtering CPU load monitoring output for privacy reasons
or system reboot only when rack power consumption $< X$

Modern Server



P²C Prototype for BMC

- P²C prototype is built on top of OpenBMC
 - Intercepts API calls and applies P²C authorization framework
- Two deployment prototypes:
 - QEMU running OpenBMC managing single server
 - Standalone daemon making RedFish calls to 8-server rack

Examples of Operations

UPPER – party
lower – party's policy

Type	Operation	Description	P ² C
Operation	Power on, off, reboot \$system	Reset \$system	(H SO and cp) or (CP and hso)
Policy			
<code>hso.pre: SUM(External.PDU.Outlets.Power) < 1000</code>			

Examples of Operations

UPPER – party
lower – party's policy

Type	Operation	Description	P ² C
Operation	Power on, off, reboot \$system	Reset \$system	(H SO and cp) or (CP and hso)
Policy			
<code>cp.pre: COUNT (External.Servers.State=="Up") >= 6</code>			

Examples of Operations

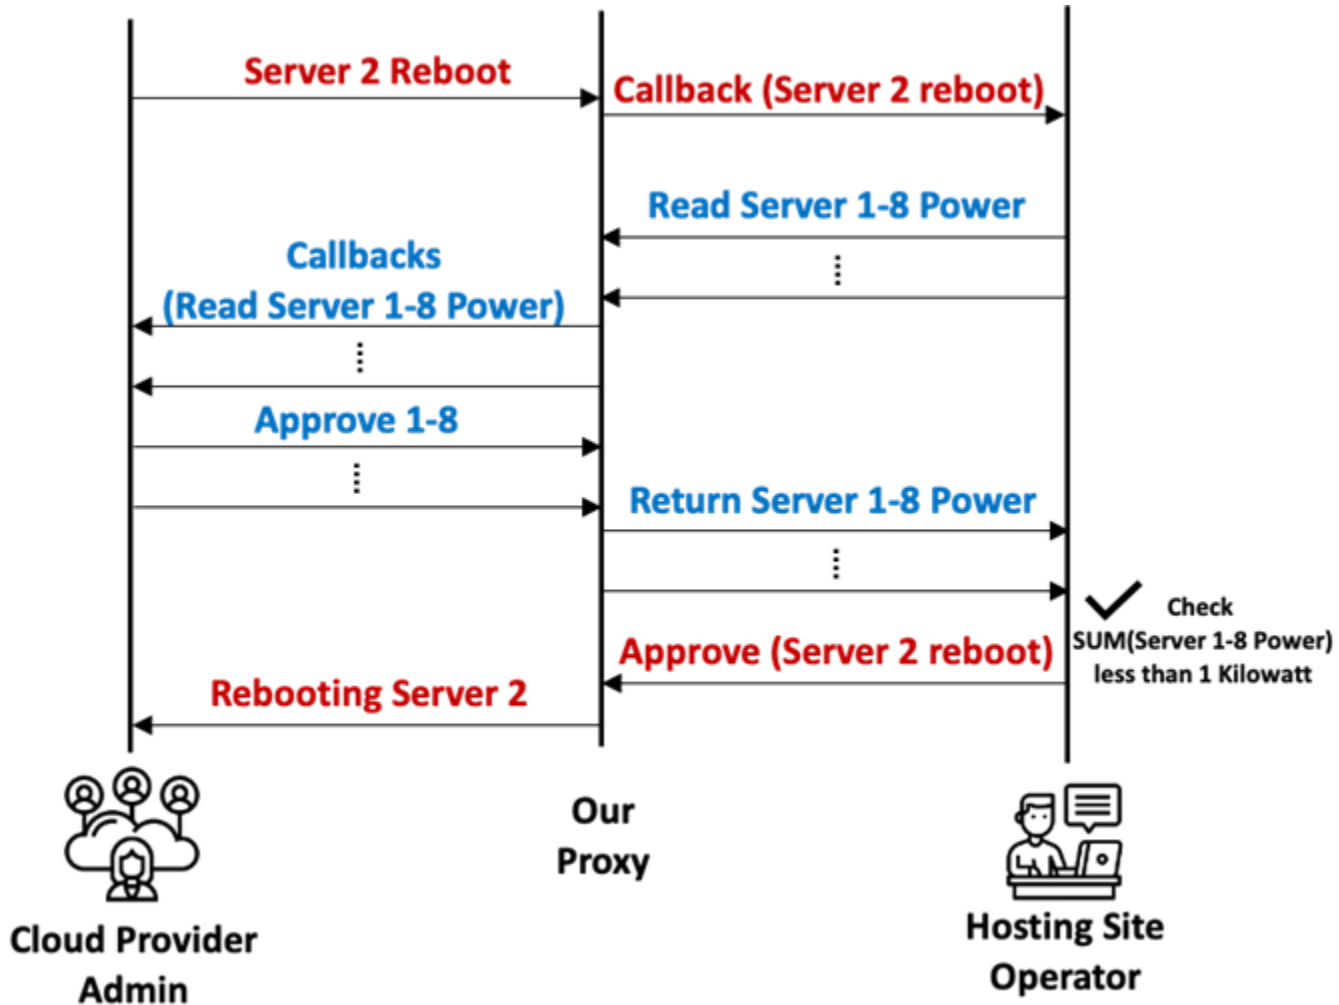
UPPER – party
lower – party's policy

Type	Operation	Description	P ² C
Logging	Get EventLog entries	Return System Event Log's collection	(H SO and cp) or (CP and hso)
Policy			
hso.post:	LogEntry.Oem*	:= null	
cp.post:	LogEntry.Severity	== "Critical"	

Examples of Operations

UPPER – party
lower – party's policy

Type	Operation	Description	P ² C
Monitoring	Get a \$chassis Chassis	Returns schema of server's \$chassis	HSO or (CP and hso)
Policy			
<code>hso.post: Chassis.Certificates := null</code>			



Open Research Directions

- Policy-as-code
 - What is the right trusted execution model for policy-as-code?
- Proxy trust model, provisioning, recovery
- Many-party control
 - Policy conflict resolving
- See paper for more
 - Networking considerations, crowdsourcing, nested P2C authorization framework, etc.

Conclusions

- Identified problem of MCloud management
- Argued that collaboration is required to succeed
- Proposed P²C framework enabling both parties to participate in MCloud management
- Described possible future research directions

Q&A

MCloud Today

AWS Outpost/Wavelength:

- Amazon builds/installs/operates server/rack
- Client provides site facility

Azure Stack Edge:

- Microsoft builds/operates server
- Client installs server and provides site facility

Azure Operator Nexus:

- Microsoft provides BoM specs
- Client builds/installs/operates whole rack

Google Anthos:

- Google provides reference specification
- System Integrator builds whole rack
- Client installs/operates whole rack