


# Operating ARC CEs: sysadmin toolbox and resource usage analysis

Accounting and sysadmin tools in ARC6



# ARCCTL: CE under your control

## Deployment Automation

- CA Certificates
- VOMS
- Firewall
- Package install

## Service Control

manage ARC services based on configuration

## Configuration Management

- overview
- validation
- in-line help
- runtime config

## Jobs Control

- get info, logs, script
- kill and clean
- ownership
- statistics

arcctl

## Accounting, Usage Analysis

- who is using CE
- how they using CE
- how much CPU, memory, disc, network

## Worker Node Environment

- features available during job runtime
- environment setup
- containers support

## Data Staging


- cache control
- transfers stats

**ARC Control Tool**  
designed to automate ARC CE operations for sysadmins



# ARCCTL Highlights

- **Single entry point** to manage ARC and related subsystems
  - one-stop-shop for sysadmins of ARC6 sites
- **Easy to use** without knowledge of ARC internals
  - quick start for new ARC CE admins
- **Shortcuts** to automate typical CE administrations and management operations
  - effective maintenance and troubleshooting
- **Modular** and **extendable**
- Designed with **bash-completion in mind**
  - Easy to find component, ID, etc
  - Completes names, Job IDs, Certificate DNs, etc



# ARCCTL: Deployment and Service Control

- Install ARC CE package\*
  - `yum install nordugrid-arc-arex`
- Deploy CA certificates
  - `arcctl deploy igtcf-ca classic`
- Fetch VOMS trust chains
  - `arcctl deploy voms-lsc -e atlas`
- Modify `arc.conf`
  - Storage areas and cache
  - Batch system and queues
  - Authorization rules
- Configure firewall
  - `arcctl deploy iptables-config`
- Run it “as-configured”
  - `arcctl service start --as-configured`



\*Available from EPEL or NorduGrid repositories



# ARCCTL: Analyze the ARC CE setup

- Read full configuration including all defaults
  - `arcctl config dump`
- Find important paths (system dirs and logs)
  - `arcctl config brief`
- Apply config tuning using inline or **online** configuration reference
  - `arcctl config describe`
- Verify configuration syntax
  - `arcctl config verify`

```
[root ~]$ arcctl config brief
ARC Storage Areas:
Control directory:
    /var/spool/arc/jobstatus
Session directories:
    /home/grid/arc6.univ/session
Scratch directory on Worker Node:
    Not configured
Cache directories:
    /home/grid/arc6.univ/cache
Additional user-defined RTE directories:
    Not configured
Accounting ar...
```

```
ARC Log Files:
A-REX Service
    /var/log/
A-REX Jobs l
    /var/log/
A-REX Helpers
    /var/log/
A-REX WS Inte
    /var/log/
A-REX Data St
    Not confi
```

**defaultttl**  
Synopsis: `defaultttl = [ttl [ttr]]`  
Description: The ttl parameter sets the time in seconds for how long a job session directory will survive after job execution has finished. If not specified the default is 1 week. The ttr parameter sets how long information about a job will be kept after the session directory is deleted. If not specified, the ttr default is one month.  
Default: `604800 2592000`  
Example:

```
[root ~]$ arcctl config describe arex defaultttl
## defaultttl = [ttl [ttr]] - The ttl parameter sets the time
## directory will survive after job execution has finished. If not specified
## the default is 1 week. The ttr parameter sets how long information about a job will be kept
## after the session directory is deleted. If not specified, the ttr default is one month.
## default: 604800 2592000
```

```
Infosys LDAP/BDII logs:
    /var/log/arc/bdii
```



# ARCCTL: Configuring Worker Node Environment

ARC CE has a feature of **dynamic and flexible provisioning** of the **execution environments on the WN (RunTimeEnvironment(RTE) concept)**.


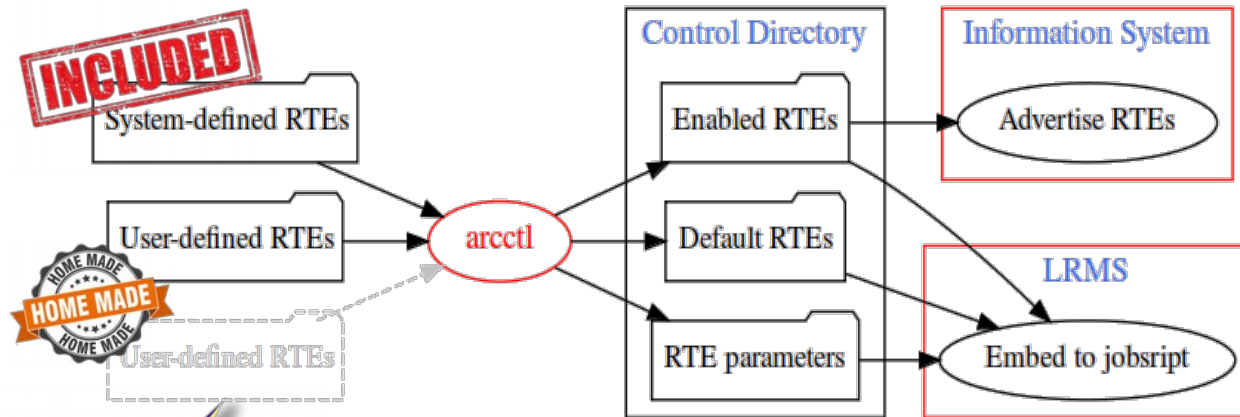
Various **RTE categories** to advertise and configure:

- **feature** (ENV/PROXY)
- **resource** (ENV/GPU)
- **software** (APPS/HEP/ATLAS)



**Out of the box RTEs are included!**

- Enable to be requested on-demand
  - `arcctl rte enable ENV/PROXY`
- Configure for all jobs
  - `arcctl rte params-set ENV/SINGULARITY SINGULARITY_IMAGE default:/srv/containers/centos7.img`
  - `arcctl rte default ENV/SINGULARITY`



# NOTICE

Only subset of command line arguments shown here!

# ARCCTL: Working with CE Jobs

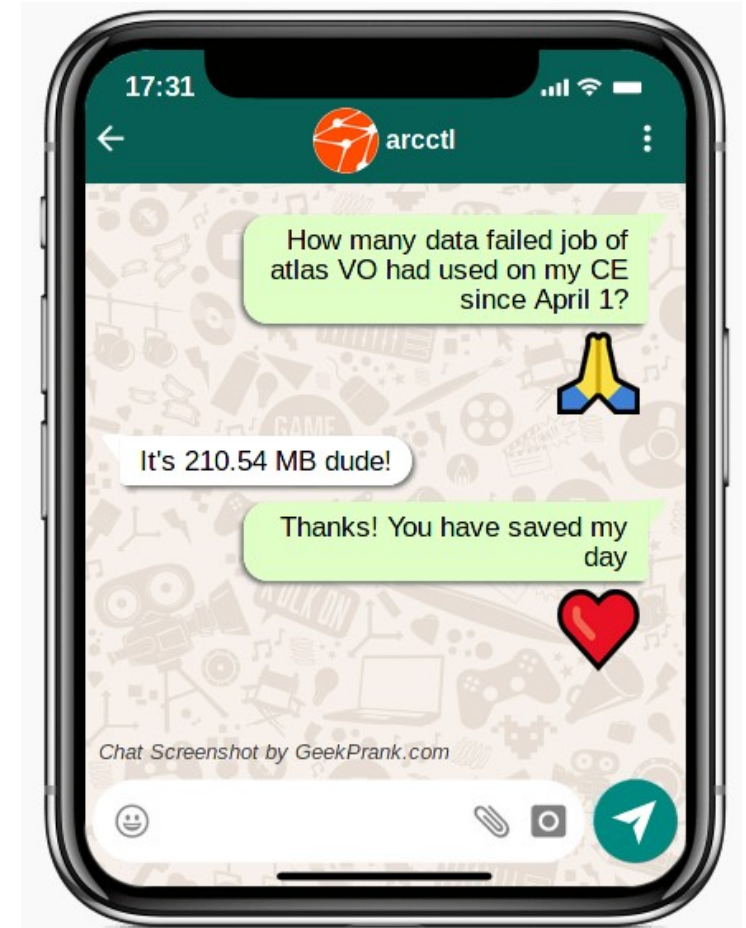
- Check jobs
  - `arcctl job list [-o OWNER]`
  - `arcctl job stats`
- Job details
  - `arcctl job info JOBID`
  - `arcctl job attr JOBID [attr]`
- Job processing logs
  - `arcctl job log JOBID [-s]`
- Job batch script and output
  - `arcctl job script JOBID`
  - `arcctl job std{out|err} JOBID`
- Job control
  - `arcctl kill JOBID`
  - `arcctl cleanall -o OWNER`

```
[root ~]$ arcctl -d INFO job killall -o "/DC=org/DC=ugrid/O=people/O=KNU/  
CN=Andrii Salnikov"  
[2019-11-14 14:32:54] [Arc.gm-jobs] [INFO] [5927/13526032] Job: 4PL0DmbAu  
lvnjw05upha610qABFKDmABFKDmENIKDmEBFKDmWik1Mm : Cancel request put  
[2019-11-14 14:32:54] [Arc.gm-jobs] [INFO] [5927/13526032] Job: LLCMDmklwE  
mvnjw05upha610qABFKDmABFKDmHHJKDmABFKDmb0EFJo : Cancel request put  
[2019-11-14 14:32:54] [Arc.gm-jobs] [INFO] [5927/13526032] Job: KJ0NDmZxB  
mvnjw05upha610qABFKDmABFKDmENIKDmHBFKd6uS52n : Cancel request put  
[2019-11-14 14:32:54] [Arc.gm-jobs] [INFO] [5927/13526032] Job: JCILDmN7s  
mvnjw05upha610qABFKDmABFKDm01IKDmABFKDmNpLeLm : Cancel request put
```



# ARCCTL: Accounting and Usage Analysis


- General accounting data queries
  - `arcctl accounting stats`
- Per-job resource usage analysis
  - `arcctl accounting job info`
- Job life-cycle events
  - `arcctl accounting job events`
- Job data transfers detailed history
  - `arcctl accounting job transfers`





# Accounting: keep track on resource usage

- ARC CE accounting subsystem collects data about **per-job resource usage** in the local **SQLite database** for **on-site CE auditing** and **publishing**
- A-REX Accounting Record\* (**AAR**) defines data stored about the job




*\*Complete AAR Definition can be found in the [documentation](#).*



# Accounting highlights: database

- Accounting DB is **integral part** of A-REX (v 6.4) and **always populated**
  - Approx. **10MB** of data per **5000** jobs
  - **SQLite**: no need to deploy additional services
  - Highly scalable



Scalability tested on 10 000 000 jobs database that was recreated from data collected for real jobs usage records archived during 5 year!



# Accounting: collecting and publishing data

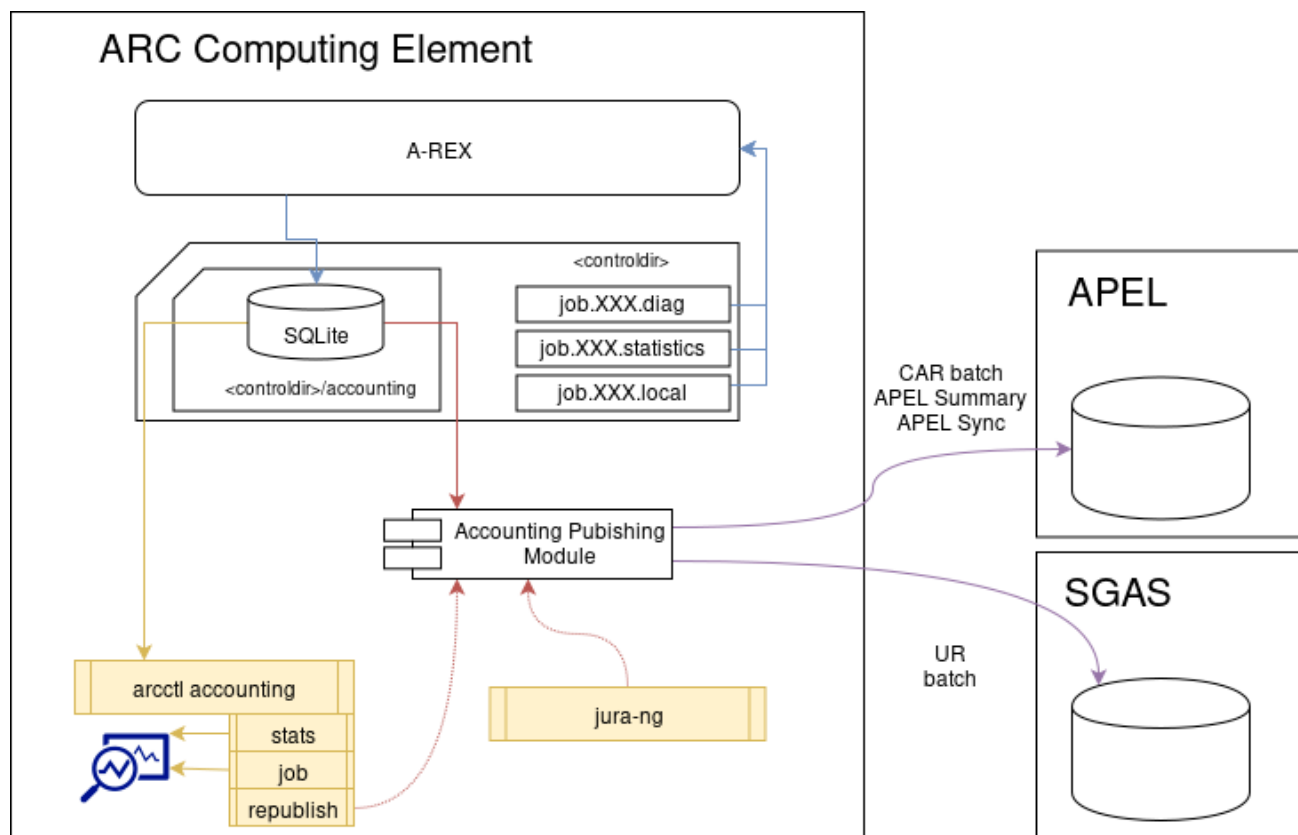
- Support for various **resource usage measurement** methods:
  - batch system data
  - GNU Time tool
  - **arc-job-cgroup** tool



- **Job Usage Reporter (JURA)** subsystem for **publishing records** to central accounting services:
  - Multiple simultaneous targets (**SGAS** and **APEL** are supported)
  - **OGF.98 UR, EMI CAR 1.2, APEL Summary** and **APEL Sync** formats
  - Highly configurable per-target



# Accounting: technical details



- **Database is a central point** of the accounting
- **Updated by A-REX** on each job state change event
- Source of data for **local analysis**
- Source of data for **publishing**
  - old data can be easily (re)published to newly added target



# Local analysis: arcctl job info

```
[root ~]# arcctl accounting job info rUyNDmQoRpvHHSrXqgZnWSqksst0oABFKDm7eKKDmAB
Job rUyNDmQoRpvHHSrXqgZnWSqksst0oABFKDm7eKKDmAB accounting info:
```

```
=====
```

## Job description:

Job was **submitted** at 2019-11-14 10:45:47 via "org.nordugrid.gridftpjob" interface using "gsiftp://atlas.bluegrass.nsc.liu.se:2811/jobs" endpoint.

Job **owned** by "/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=atlact1/CN=Robot: ATLAS aCT 1" as a member of "atlas" **WLCG VO**.

It was targeted to the "bluegrass" queue with "1906417" **LRMS ID**.

Job completed with **exit code 0** at 2019-11-14 12:44:14.

Following job properties are recorded:

**Localuser:** atlasuser0000:atlasuser0000

**Lrms:** SLURM

**Systemsoftware:** CentOS Linux 7 (Core)

**Nodenames:** n1535

**Benchmark:** HEPSPEC:15.925

**Jobname:** user.gfrattar.410470.PhPy8EG\_ttbar\_FTAG1.pflow.mc16e.1Nov/.4546575012



# Local analysis: arcctl job info (continued)

## Resource usage:

Execution timeframe: 2019-11-14 10:45:47 - 2019-11-14 12:44:14

Used WallTime: 2296

Used CPUTime: 2358 (including 62 of kernel time)

Used WN Scratch: 21.3M

Max physical memory: 1.3M

Max virtual memory: 1.3M

Used CPUs: 1 on 1 node(s)

Data staged in: 585.2M

Data staged out: 7.1M

## Used RunTime Environments:

ENV/PROXY

APPS/HEP/ATLAS-SITE

ENV/RTE

## Auth token attributes provided:

VOMS FQAN: /atlas/Role=pilot

VOMS FQAN: /atlas

VOMS FQAN: /atlas/lcg1



```
[root ~]# arcctl accounting job info aa4LdM5m3Ivnjw05upha610qABFKDmABF
KDm5FKDmPqGKdMf2jwun -o json
{"RunTimeEnvironments": [], "UserSN": "/DC=EU/DC=EGI/C=HR/O=Robots/O=S
RCE/CN=Robot:argo-egi@cro-ngi.hr", "NodeCount": 1, "VOID": 3, "JobEven
ts": [{"LRMSSTART": "2019-08-17 05:37:04"}, {"LRMSEND": "2019-08-17 05
:37:05"}, {"ACCEPTED": "2019-08-17 05:37:52"}, {"PREPARING": "2019-08-
17 05:37:52"}, {"DTRDOWNLOADSTART": "2019-08-17 05:37:52"}, {"DTRDOWNL
OADEND": "2019-08-17 05:37:52"}, {"SUBMIT": "2019-08-17 05:37:53"}, {"
INLRMS": "2019-08-17 05:37:53"}, {"FINISHING": "2019-08-17 05:42:20"},
{"DTRUPLOADSTART": "2019-08-17 05:42:20"}, {"DTRUPLOADEND": "2019-08-
17 05:42:33"}, {"FINISHED": "2019-08-17 05:42:34"}], "UserID": 2, "Job
ID": "aa4LdM5m3Ivnjw05upha610qABFKDmABFKDm5FKDmPqGKdMf2jwun", "Queue"
: "grid_rt", "WLCGVO": "ops", "StageOutVolume": 24, "CPUTime": 1, "En
dTime": "2019-08-17 02:42:34", "UsedScratch": 746, "UsedVirtMem": 1056
, "DataTransfers": [{"url": "srm://se05.esc.qmul.ac.uk/ops/nagios-argo
-mon.egi.cro-ngi.hr/arcce/srm-input", "timestart": "2019-08-17 02:37:5
2", "type": "cache_input", "timeend": "2019-08-17 02:37:52", "size": 5
7}, {"url": "srm://se05.esc.qmul.ac.uk/ops/nagios-argo-mon.egi.cro-ngi
.hr/arcce/srm-20190817T0437-arc6.univ.kiev.ua", "timestart": "2019-08-
17 02:42:20", "type": "output", "timeend": "2019-08-17 02:42:33", "siz
e": 24}], "Status": "completed", "QueueID": 3, "JobExtraInfo": {"local
user": "ops23:ops", "jobname": "org.nordugrid.ARC-CE-SRM-result-ops",
"systemsoftware": "CentOS Linux 7 (Core)", "clienthost": "161.53.0.239
:55218", "lrms": "pbs", "nodenames": "s3"}, "AuthTokenAttributes": [{"
mainfqan": "/ops"}], "RecordID": 1832, "UsedCPUTime": 0, "EndpointURL"
: "https://arc6.univ.kiev.ua:443/arex", "UsedMemory": 1056, "UsedWallT
ime": 1, "SubmitTime": "2019-08-17 02:37:52", "StageInVolume": 0, "End
pointID": 1, "Interface": "org.ogf.glue.emies.activitycreation", "Used
CPUUserTime": 0, "ExitCode": 0, "StatusID": 3, "UsedCPUKernelTime": 0,
"LocalJobID": "2817764"}
```

Want it machine-readable?

Use "-o json"!



# ARCCTL: analysis of job events and transfers

```
[root ~]# arcctl accounting job events  
rUyNDmQoRpvNHHSrXqgZnWSqkstOoABFKDm7eKKDmAB
```

```
2019-11-14 11:45:47 ACCEPTED  
2019-11-14 11:45:47 PREPARING  
2019-11-14 11:45:47 DTRDOWNLOADSTART  
2019-11-14 13:02:38 DTRDOWNLOADEND  
2019-11-14 13:02:39 SUBMIT  
2019-11-14 13:02:52 INLRMS  
2019-11-14 13:03:26 LRMSSTART  
2019-11-14 13:41:42 LRMSSEND  
2019-11-14 13:42:50 FINISHING  
2019-11-14 13:42:53 DTRUPLOADSTART  
2019-11-14 13:44:13 DTRUPLOADEND  
2019-11-14 13:44:14 FINISHED
```

```
[root ~]# arcctl accounting job transfers  
rUyNDmQoRpvNHHSrXqgZnWSqkstOoABFKDm7eKKDmAB
```

```
Data transfers (downloads) performed during A-REX stage-in:  
http://pandaserver.cern.ch:25085/cache/pilot/pilot2.tar.gz
```

```
(from cache):
```

```
Size: 258.4K
```

```
Download timeframe: 2019-11-14 10:45:47 - 2019-11-14
```

```
10:45:50
```

```
rucio://rucio-lb-prod.cern.ch/replicas/mc1\_13TeV/
```

```
DAOD\_FTAG1.19607426.\_001748.pool.root.1:
```

```
Size: 564.7M
```

```
Download timeframe: 2019-11-14 10:45:48 - 2019-11-14
```

```
11:58:51
```

```
<output omitted>
```

```
Data transfers (uploads) performed during A-REX stage-out:
```

```
srm://srm.ndgf.org:8443/srm/managerv2?SFN=/atlas/disk/
```

```
atlasscratchdisk/rucio/panda/42/82/
```

```
panda.um.user.gfrattar.19770512.\_001292.output.root:checksumty
```

```
pe=adler32:checksumvalue=75aa3e91:
```

```
Size: 7.1M
```

```
Upload timeframe: 2019-11-14 12:42:53 - 2019-11-14 12:44:13
```

```
<output omitted>
```



# ARCCTL: Resource usage statistics

```
[root arc-dev]# arcctl accounting stats
```

```
A-REX Accounting Statistics:
```

```
Number of Jobs: 238
```

```
Execution timeframe: 2019-07-09 18:42:18 - 2019-07-12 15:16:56
```

```
Total WallTime: 6 days, 1:35:37
```

```
Total CPUTime: 4:19:45 (including 0:00:00 of kernel time)
```

```
Data staged in: 62.1M
```

```
Data staged out: 1.7K
```

```
[root arc-dev]# arcctl accounting stats --filter-state failed --filter-vo atlas \  
> --start-from 2019-04-01 -o data-staged-in  
220767192
```

```
[root arc-dev]# arcctl accounting stats --filter-queue grid \  
> --filter-user /DC=\org/DC=\ugrid/O=\people/O=\KNU/CN=\Andrii\  
A-REX Accounting Statistics:
```

```
Number of Jobs: 8
```

```
Execution timeframe: 2019-07-09 18:42:18 - 2019-07-10 17:31:59
```

```
Total WallTime: 0:01:20
```

```
Total CPUTime: 0:00:00 (including
```

```
Data staged in: 61.7M
```

```
Data staged out: 0
```

```
[root arc-dev]# arcctl accounting stats --filter-extra jobname "test 04" -o jobids  
g3AMDmCCP5unf5481mks8bjnABFKDmABFKDmN9IKDmIBFKDmU0ym  
nEhMDmvqQ5unf5481mks8bjnABFKDmABFKDmN9IKDmZBFKDtXEYGn  
UufLDmmsS5unf5481mks8bjnABFKDmABFKDmSMKDmNBFKDm9nRnVo  
bdoLDmCnS5unf5481mks8bjnABFKDmABFKDmSMKDmABFKDmfKldDn
```

```
[root arc-dev]# arcctl accounting stats --filter-extra rte ENV/PROXY -o json | jq .
```

```
{  
  "stageout": 0,  
  "rangeend": 1562938441,  
  "count": 86,  
  "cpukerneltime": 0,  
  "users": [  
    "/DC=EU/DC=EGI/C=HR/O=Robots/O=SRCE/CN=Robot:argo-egi@cro-ngi.hr",  
    "/DC=EU/DC=EGI/C=GR/O=Robots/O=Greek Research and Technology Network/CN=Robot:argo-egi@grnet.gr"  
  ],  
  "wlcgvos": [  
    "ops"  
  ],  
  "cpusertime": 0,  
  "cputime": 0,  
  "rangestart": 1562688395,  
  "stagein": 398180,  
  "walltime": 172  
}
```





# WARNING: ARCCTL is highly addictive!

- **Once upon a time** there was a sysadmin who needed to **install CA certificates** and **VOMS trust chains** on a **RUCIO machine**.
- He hesitated a bit to install them manually...
- ...but end up **installing A-REX only to get access to ARCCTL!**
- **Therapy:** we'll provide a **standalone arcctl** in the future ARC releases for admins get hooked on it.



# Thank you for your kind attention!

## Questions?



One Tool to rule them all,  
One Tool to find them,  
One Tool to bring them all  
and in the darkness bind them.

