

# Triton

Triton und Manta (Private) Cloud auf Basis von SmartOS

Töns Büker

fraosug - Jitsi, 15.09.2020

# Joyent - Hosting und Cloud

- ▶ **Unternehmen** Gegründet 2004/2005
- ▶ **Projekte** Node.js, SmartOS, Illumos, pkgsrc
- ▶ **OpenSource** SmartOS - CDDL (2011), SmartDataCenter und Manta - MPL (2014)
- ▶ **Triton** Namenswechsel (2015)
- ▶ **Samsung** Übernahme - Teil der Mobilstrategie (Juni 2016)
- ▶ **Joyent Public Cloud EOL** Kein Geschäft mehr für Samsung (November 2019)

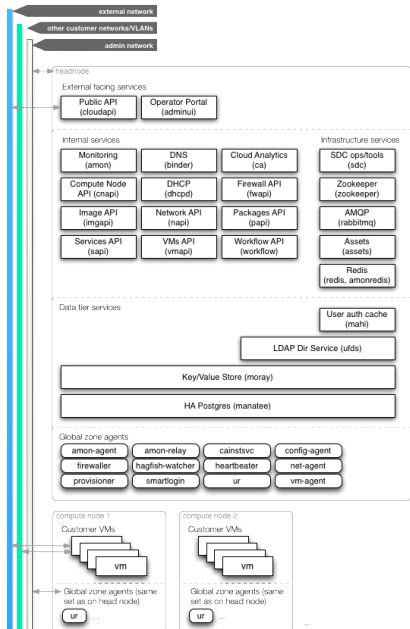


Joyent

Woraus besteht Triton?

Woraus besteht Manta?

# Triton Dienste - Überblick



# Triton Dienste - im Einzelnen

- ▶ **Manatee, Moray** “Source of Truth” - automatisierter Postgres-Cluster mit Key/Value Store
- ▶ **Interne Dienste** diverse APIs, Monitoring, DNS, DHCP, portolan, etc.
- ▶ **Externe Dienste** Cloud API, Docker API, Operations Portal
- ▶ **CNS** Container Naming Service
- ▶ **CMon** Container Monitoring Service
- ▶ **VolAPI** Shared Volumes API
- ▶ **prometheus** Prometheus für Triton/Manta-Dienste
- ▶ **grafana** Grafana für Triton/Manta-Dienste
- ▶ **diverse Agenten** auf den Compute Nodes



# Triton Installation

- ▶ **CoaL** Cloud on a Laptop
- ▶ **Netzwerk, Domain** Angaben zu VLANs, IP-Netzen und DNS Domain machen
- ▶ **Erster Login** Tipp: Nicht vergessen auf Channel “release” um- und zfs compression einzustellen
- ▶ **Post-Setup** DC-Verlinkung, weitere CNs pxe-booten, HA-Manatee ausrollen, weitere Dienste ausrollen
- ▶ **Updates** auf Imagebasis (PI, Zoneimages, Agenten) - [gist.github.com/bahamat/](https://gist.github.com/bahamat/)
- ▶ **Images** [images.joyent.com](https://images.joyent.com) - Quelle für OS- und Applikationsimages
- ▶ **Packages** Beispielpakete sind ebenfalls enthalten

# Triton Highlights I

- ▶ **bhyve** neue Virtualisierung in SmartOS/Triton
- ▶ **sdc-migrate** Offline Migrieren von VMs und Zonen
- ▶ **triton-dehydrated** letsencrypt Zertifikate für Instanzen und API-Endpoints
- ▶ **Linux CN** Ein Linux ComputeNode für Linux Container Workloads
- ▶ **triton CLI** network, volumes, etc.
- ▶ **hermes/logarchiver** Zentrales Logmanagement
- ▶ **Fabric Networks** virtuelle Netzwerke (SDN)

## Triton Highlights II

- ▶ **Linking Triton Datacenters** User Federation, SSO-Service?
- ▶ **sdc-vmadm** Welche Instanz läuft wo?
- ▶ **kmdb / Crypted Data at Rest** Verschlüsselung, API, Keys, Zufallszahlen
- ▶ **Terraform Provider** auch hier Volumes und delegated-datasets
- ▶ **Portal?** github-Repos mit sdc-piranha gesichtet - und anderen...
- ▶ **piadm** bootfähiger ZFS-Pool - zielt auf Bare Metal Hoster
- ▶ **RFDs** Kommende Projekte ([github.com/joyent/rfd](https://github.com/joyent/rfd))



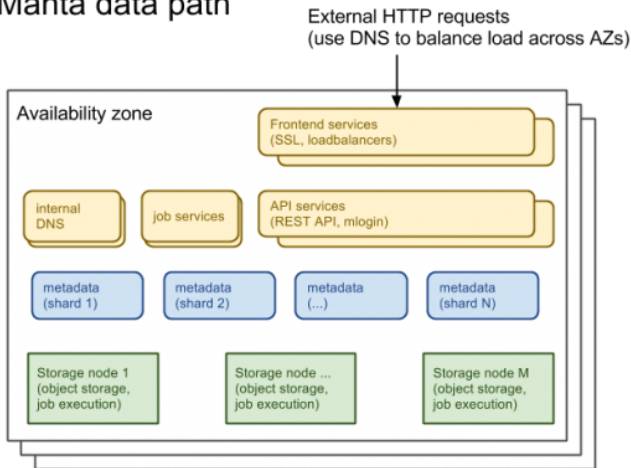
Joyent

Woraus besteht Triton?

Woraus besteht Manta?

# Manta Überblick

## Manta data path



■ Stateless services

■ Metadata tier

■ Storage tier

- Stateless services use internal DNS to balance load across all AZs
- Metadata shards have at least one instance in each AZ
- Copies of objects are stored in separate AZs (up to three)

# Manta Dienste

- ▶ **Object Storage** HTTP (REST), signed URLs, Unterverzeichnisse, multi-part Uploads
- ▶ **Jobs** “Compute on Data” - “Map and Reduce”
- ▶ **SnapLinks** Versionierung
- ▶ **Docs** <https://apidocs.joyent.com/manta/>

# Manta Monitoring - MadTom

## ▶ MadTom Dienstübersicht - Monitoring Microservices

Host Type	🏠		🏠		🏠								
loadbalancer	896353b3 stud: <span style="color: green;">■</span> <span style="color: green;">■</span> haproxy: <span style="color: green;">■</span> <span style="color: green;">■</span>	7b7b679d stud: <span style="color: green;">■</span> <span style="color: green;">■</span> haproxy: <span style="color: green;">■</span> <span style="color: green;">■</span>	c9504460 stud: <span style="color: green;">■</span> <span style="color: green;">■</span> haproxy: <span style="color: green;">■</span> <span style="color: green;">■</span>	3e682888 stud: <span style="color: green;">■</span> <span style="color: green;">■</span> haproxy: <span style="color: green;">■</span> <span style="color: green;">■</span>	979a1712 stud: <span style="color: green;">■</span> <span style="color: green;">■</span> haproxy: <span style="color: green;">■</span> <span style="color: green;">■</span>	0359f512 stud: <span style="color: green;">■</span> <span style="color: green;">■</span> haproxy: <span style="color: green;">■</span> <span style="color: green;">■</span>							
webapi	6bcf530e haproxy: <span style="color: green;">■</span> <span style="color: green;">■</span> muskie: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>	82c22bfb haproxy: <span style="color: green;">■</span> <span style="color: green;">■</span> muskie: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>	b7727a32 haproxy: <span style="color: green;">■</span> <span style="color: green;">■</span> muskie: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>	ba632985 haproxy: <span style="color: green;">■</span> <span style="color: green;">■</span> muskie: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>	80084ec9 haproxy: <span style="color: green;">■</span> <span style="color: green;">■</span> muskie: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>	38d0a4e6 haproxy: <span style="color: green;">■</span> <span style="color: green;">■</span> muskie: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>							
authcache	c77515b2 mahi-ser: <span style="color: green;">■</span> <span style="color: green;">■</span> mahi-rep: <span style="color: green;">■</span> <span style="color: green;">■</span> redis: <span style="color: green;">■</span> <span style="color: green;">■</span>		7fda0dcb mahi-ser: <span style="color: green;">■</span> <span style="color: green;">■</span> mahi-rep: <span style="color: green;">■</span> <span style="color: green;">■</span> redis: <span style="color: green;">■</span> <span style="color: green;">■</span>										
electric-moray	50a5c8e1 haproxy-tcp: <span style="color: green;">■</span> <span style="color: green;">■</span> electric-moray: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>	bdc1e02e haproxy-tcp: <span style="color: green;">■</span> <span style="color: green;">■</span> electric-moray: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>	a81f137f haproxy-tcp: <span style="color: green;">■</span> <span style="color: green;">■</span> electric-moray: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>	2e4f09e2 haproxy-tcp: <span style="color: green;">■</span> <span style="color: green;">■</span> electric-moray: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>	5d734811 haproxy-tcp: <span style="color: green;">■</span> <span style="color: green;">■</span> electric-moray: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>	ca28c843 haproxy-tcp: <span style="color: green;">■</span> <span style="color: green;">■</span> electric-moray: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>							
moray	4e89d72 haproxy-tcp: <span style="color: green;">■</span> <span style="color: green;">■</span> moray: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>	439ac915 haproxy-tcp: <span style="color: green;">■</span> <span style="color: green;">■</span> moray: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>	b2ca7086 haproxy-tcp: <span style="color: green;">■</span> <span style="color: green;">■</span> moray: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>	a764b57b haproxy-tcp: <span style="color: green;">■</span> <span style="color: green;">■</span> moray: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>	f544a4 haproxy-tcp: <span style="color: green;">■</span> <span style="color: green;">■</span> moray: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>	cd8a7d32 haproxy-tcp: <span style="color: green;">■</span> <span style="color: green;">■</span> moray: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>							
postgres	24f89ce1 manatee: <span style="color: green;">■</span> <span style="color: green;">■</span> postgres: <span style="color: green;">■</span> <span style="color: green;">■</span> pg-backup: <span style="color: green;">■</span> <span style="color: green;">■</span>	8961a64c manatee: <span style="color: green;">■</span> <span style="color: green;">■</span> postgres: <span style="color: green;">■</span> <span style="color: green;">■</span> pg-backup: <span style="color: green;">■</span> <span style="color: green;">■</span>	bfa1f357 manatee: <span style="color: green;">■</span> <span style="color: green;">■</span> postgres: <span style="color: green;">■</span> <span style="color: green;">■</span> pg-backup: <span style="color: green;">■</span> <span style="color: green;">■</span>	cab1a03d manatee: <span style="color: green;">■</span> <span style="color: green;">■</span> postgres: <span style="color: green;">■</span> <span style="color: green;">■</span> pg-backup: <span style="color: green;">■</span> <span style="color: green;">■</span>	9e1f0898 manatee: <span style="color: green;">■</span> <span style="color: green;">■</span> postgres: <span style="color: green;">■</span> <span style="color: green;">■</span> pg-backup: <span style="color: green;">■</span> <span style="color: green;">■</span>	323ca13b manatee: <span style="color: green;">■</span> <span style="color: green;">■</span> postgres: <span style="color: green;">■</span> <span style="color: green;">■</span> pg-backup: <span style="color: green;">■</span> <span style="color: green;">■</span>							
storage	aa28635e mako: <span style="color: green;">■</span> <span style="color: green;">■</span> minnow: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>	34b313da mako: <span style="color: green;">■</span> <span style="color: green;">■</span> minnow: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>	0a55ab63 mako: <span style="color: green;">■</span> <span style="color: green;">■</span> minnow: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>	823d5091 mako: <span style="color: green;">■</span> <span style="color: green;">■</span> minnow: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>	267e3fab mako: <span style="color: green;">■</span> <span style="color: green;">■</span> minnow: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>	095abf01 mako: <span style="color: green;">■</span> <span style="color: green;">■</span> minnow: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>	fb2bf521 mako: <span style="color: green;">■</span> <span style="color: green;">■</span> minnow: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>	b790f563 mako: <span style="color: red;">■</span> <span style="color: red;">■</span> minnow: <span style="color: red;">■</span> <span style="color: red;">■</span> <span style="color: red;">■</span> <span style="color: red;">■</span>	000d5e9 mako: <span style="color: green;">■</span> <span style="color: green;">■</span> minnow: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>	6bc0f8bc mako: <span style="color: green;">■</span> <span style="color: green;">■</span> minnow: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>	7c31987e mako: <span style="color: green;">■</span> <span style="color: green;">■</span> minnow: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>	5e93487a mako: <span style="color: green;">■</span> <span style="color: green;">■</span> minnow: <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span> <span style="color: green;">■</span>	
agent	5f3dcb8f agent: <span style="color: green;">■</span> <span style="color: green;">■</span>	506ba2f5 agent: <span style="color: green;">■</span> <span style="color: green;">■</span>	0e85c0b8 agent: <span style="color: green;">■</span> <span style="color: green;">■</span>	bds9c598 agent: <span style="color: green;">■</span> <span style="color: green;">■</span>	83123a17 agent: <span style="color: green;">■</span> <span style="color: green;">■</span>	1600a3b0 agent: <span style="color: green;">■</span> <span style="color: green;">■</span>	98b2cb03 agent: <span style="color: green;">■</span> <span style="color: green;">■</span>	71c27a4f agent: <span style="color: red;">■</span> <span style="color: red;">■</span>	00000000 agent: <span style="color: green;">■</span> <span style="color: green;">■</span>	8a0021e4 agent: <span style="color: green;">■</span> <span style="color: green;">■</span>	697e7f5f agent: <span style="color: green;">■</span> <span style="color: green;">■</span>	bca7e918 agent: <span style="color: green;">■</span> <span style="color: green;">■</span>	
jobsupervisor					824cb40 supervisor: <span style="color: green;">■</span> <span style="color: green;">■</span>					1afdf1e supervisor: <span style="color: green;">■</span> <span style="color: green;">■</span>			
nameservice					15cbw680 binder: <span style="color: green;">■</span> <span style="color: green;">■</span> zookeeper: <span style="color: green;">■</span> <span style="color: green;">■</span>					fd41cw75 binder: <span style="color: green;">■</span> <span style="color: green;">■</span> zookeeper: <span style="color: green;">■</span> <span style="color: green;">■</span>			
medusa					513cb09c medusa: <span style="color: green;">■</span> <span style="color: green;">■</span>								

# Manta Monitoring - Marlin Dashboard

## ▶ Marlin Jobs und Marlinzonen

Manta Marlin Empty Dashboard updated: 10:18:07

Warning: first of 1 error; connect timeout after 5000ms

### Agents

DC Host	Started	Disk stop used mem stop used	Tasks	Z	B	R	D	Zones
10-12T16:19:15	0 / 340B 0 / 39305MB	4 47 3 1 0						
10-11T20:19:14	0 / 314GB 0 / 78617MB	7 95 0 0 0						
10-13T06:19:17	0 / 10GB 0 / 22931MB	3 31 0 0 0						
10-11T11:19:06	0 / 10GB 0 / 26207MB	13 31 0 0 0						
10-12T21:19:05	0 / 185GB 0 / 52402MB	5 64 1 0 0						
10-12T21:19:05	0 / 63GB 0 / 39305MB	5 47 3 1 0						
10-13T06:19:17	0 / 610B 0 / 34390MB	8 41 1 0 0						
10-09T12:19:05	0 / 20GB 0 / 39305MB	2 47 1 0 0						
10-12T18:19:14	120 / 232GB 1024 / 78617MB	12 98 1 2 0						
10-12T08:19:19	0 / 4GB 0 / 52403MB	6 63 1 0 0						
10-09T02:19:15	0 / 4GB 0 / 52403MB	2 64 1 0 0						

□ Idle (ready to run user task) 
 ■ Busy (running user task) 
 ■ Resetting 
 ■ Disabled (requires operator intervention)

### Active jobs

JobId	User	Started	Z	A	lp	Td	Tc	O	R	Err
73959b57-ae7b-ed9c-aa3-a73ebac7a66	poseidon	08-11T17:19:09	0	25	6	18	18	0	0	13
e8a39851-4b1f-4c27-c0f1-a5e41e2cb810	poseidon	10-13T08:15:04	1	1	1	65	0	0	0	0
3d2e7830-637f-eb8d-e065-fd5d88de8a1	poseidon	10-13T08:17:03	1	1	4	16	14	0	0	0

Z = # of zones in use    A = # of supervisor assignments    lp = # of inputs processed already    Td = # of tasks dispatched

# Manta Zukunft

- ▶ **Manta-v1 vs. Manta-v2**
- ▶ **Effizienzsteigerung**

# Kontakt

- ▶ **E-Mail:** [toens.bueker@plusserver.com](mailto:toens.bueker@plusserver.com)