



## Elon Musk's Visionary Talk on the Colonization of Mars

Presented September 29, 2016 at the 67th International Astronautical Congress in Guadalajara, Mexico

Thank you very much for having me, look forward to talking about the SpaceX Mars architecture.

And what I really want to try to achieve here is to make Mars seem possible, make it seem as though it's something that we can do in our lifetimes — and that you can go, and is there really a way that anyone can go if they wanted to? I think that's really the important thing.

So first of all, why go anywhere, right? The... I think there are really two fundamental paths. History is going to bifurcate along two directions: One path is we stay on Earth forever, and then there will be some eventual extinction event — I don't have an immediate doomsday prophecy — but there's... it's eventually, history just that there will be some doomsday event.

The alternative is to become a space-faring civilization and a multi-planet species, which I hope you agree that is the right way to go.


Yes?

That's what we want.

So how do we figure out how to take you to Mars and create a self-sustaining city? A city that it is not merely an outpost, but could become planet and its own right, and thus we could become a truly multi-planet species.

Sometimes people wonder, what about other places in the solar system, why Mars?


Well, just to sort of put things in perspective, this is what, this is an actual scale of what the solar system looks like. So we're currently in the third little rock from the left.



That's Earth.

Yeah, exactly.

And our goal is to go to the fourth rock on the left — that's Mars. But you can get a sense for the real scale of the solar system, how big the Sun is, Jupiter, Neptune, Saturn, Uranus, and then the little guys and on the right are Pluto and friends.



This sort of helps see it, not quite to scale, but it gives you a better sense for where things are. So our options for going to... For becoming a multi-planet species within our solar system, are limited.

We have, in terms of nearby options, we've got Venus. But Venus is a high pressure... A super-high-pressure hot acid bath. So that would be a tricky one. Venus is not at all like the goddess. This is not, in no way similar to, to the actual goddess. So it'd be really difficult to make things work on Venus. Mercury is also way too close to the Sun. We could go potentially onto one of the moons of Jupiter, or Saturn, but those are quite far out — much further from the Sun. A lot harder to get to.

It really leaves us with one option if we want to become a multi-planet civilization, and that's Mars.

We could conceivably go to the moon, and I have nothing against going to the moon, but I think it's challenging to create a... Become multi-planetary on the moon because it's much smaller than a planet. It doesn't have any atmosphere, it's not as resource-rich as Mars, it's got a 28-day day — whereas the Mars day is 24-and-a-half hours. And in general Mars is far better suited to ultimately scale up to be a self-sustaining civilization.

Just to give some comparison between the two planets... They're actually remarkably close in a lot of ways. And in fact we now believe that early Mars was a lot like Earth. And in fact if we could warm Mars up, we would once again have a thick atmosphere and liquid oceans.

	Earth	Mars
Diameter	12,756 km	6,792 km
Average distance from sun	150,000 km	229,000 km
Temperature Range	-88° C to 58° C	-140° C to 30° C
Atmospheric Composition	78% N <sub>2</sub> , 21% O <sub>2</sub> , 1% Other	96% CO <sub>2</sub> , <2% Ar, >2% Other
Relative gravity	1.0	0.37
Day Length	24 hours	24 hours 40 minutes
Land Mass	148.9 million km <sup>2</sup>	144.8 million km <sup>2</sup>
Population	7 billion	0

So where things are right now, Mars is about half again as far from the sun as Earth. So still decent sunlight. It's a little cold, but we can warm it up. And it has a very helpful atmosphere which, in the case of Mars, being primarily CO<sub>2</sub> with some nitrogen, and argon, and few other trace elements, means that we can grow plants on Mars just by compressing the atmosphere. And it has nitrogen, too, which is very important for growing plants.

It would be quite fun because you have gravity, which is about 37% that of Earth, so you'd be able to lift heavy things and bound around and have a lot of fun. And the day is remarkably

close to that of Earth, and so we just need to change that bottom row, because currently we have 7 billion people on earth and zero on Mars.

So there's been a lot of great work by NASA and other organizations in early exploration of Mars and understanding what Mars is like, where could we land, what's the composition of the atmosphere, where is there water — water ice, I should say — and so we need to go from these early exploration missions to actually building a city.

The issue that we have today is that if you look at a Venn diagram, there's no intersection of sets of people who want to go and can afford to go. In fact, right now, you cannot go to Mars for infinite money.



Using traditional methods, you know if you've taken a sort of Apollo-style approach, an optimistic class number would be about \$10 billion a person. So for example, in the Apollo program, the cost estimates are somewhere between \$100 to \$200 billion in current-year dollars. And we sent 12 people to the surface of moon. Which was an incredible thing and probably one of the greatest achievements of humanity.

But that's a steep price to pay for a ticket. That's why these circles only just barely touch. So you can't create a self-sustaining civilization if the ticket price is \$10 billion per person.



What we need is to move those circles together.

If we can get the cost of moving to Mars to be roughly equivalent to a median house price in the US, which is around \$200,000, then I think the probability of establishing a self-sustaining civilization is very high. I think it would almost certainly occur. Not everyone would want to go — in fact, I think a relatively small number of people from Earth would want to go — but enough would want to go, and who could afford the trip, that it would happen.

You keep looking at sponsorship and I think it gets the point where almost anyone if they saved up, and this was their goal, they could ultimately save enough money to buy a ticket and move to Mars. And Mars would have a labor shortage for a long time, so jobs would not be in short supply.

So it is a bit tricky because we have to figure out how to improve the cost of trips to Mars by 5,000,000%. So this is just not easy and... I mean, it sounds like virtually impossible, but I think there are ways through it.

This translates to an improvement of approximately four-and-a-half orders of magnitude.


Full Re-usability  
Refilling in Orbit  
Propellant Production on Mars  
Right Propellant

These are the key elements that are needed in order to achieve a four-and-a-half order of magnitude improvement. Most of the improvement would come from full reusability, somewhere between two and two-and-a-half orders of magnitude. And then the other two orders of magnitude would come from refilling in orbit, propellant production on Mars, and choosing the right propellant.

So I'm going to go into detail on all those.

Full reusability is really the super-hard one.

It's very difficult to achieve reusability for even an orbital system. And that challenge becomes substantially greater for a system that has to go to another planet.



Boeing 737	
Price	\$90M
Passenger Capability	180 people
Cost/Person - Single Use	\$500,000
Cost/Person - Reusable	\$43 (LA to Las Vegas)
Cost of Fuel / Person	\$10

But as an example of the difference between reusability and expendability in aircraft — and you could actually use any form of transport, you could say a car, bicycle, horse — if they were single-use, almost no one would use them. They'd be too expensive.

But with frequent flights you can take something like an aircraft that costs \$90 million, and if it were single-use, you'd have to pay half a million dollars per flight. But you can actually buy a ticket on Southwest, right now, from LA to Vegas, for \$43 — including taxes. So that's, I mean, that's a massive improvement right there. It's showing a four-order-of-magnitude improvement.

Now this is harder. The reusability doesn't apply quite as much to Mars, because the number of times that you could reuse the spaceship — the spaceship part of the system — is less often because the Earth-Mars rendezvous only occurs every 26 months.

So you get to use the spaceship part roughly every 2 years. Now you get to use the booster and the tanker as frequently as you'd like, and that's why it really makes a lot of sense to load the spaceship into orbit, with essentially tanks dry, have it have really quite big tanks that you then use the booster and tanker to refill while it's in orbit, and maximize the payload of the spaceship, so that when it goes to Mars you really have a very large payload capability.

So refilling in orbit is one of the essential elements of this.

Not refilling in orbit would require a  
3-stage vehicle at 5-10x the size and cost

Spreading the required lift capacity across  
multiple launches substantially reduces  
Development costs and compresses schedule

Combined with re-usability, refilling makes  
performance shortfalls an incremental rather  
than an exponential cost increase

Without refilling in orbit, you would have a half-order of magnitude impact, roughly, on the cost. By half order of magnitude — I think this audience mostly knows — but what that means is, each order of magnitude is a factor of 10. So not refilling in orbit would mean a 500%, roughly, increase in the cost per ticket.

It also allows us to build a smaller vehicle and lower the development costs, although this is quite big, but it would be much harder to build something that's five to 10 times the size. And it also reduces the sensitivity of performance characteristics of the booster rocket and tanker.

So if there's a shortfall in the performance of any of the elements, you can actually make up for it by having one or two extra refilling trips to the spaceship. So this is very important for reducing the susceptibility of the system to a performance shortfall.

And then producing propellant on Mars is actually very obviously important.

Again if, if we didn't do this, it would have at least a half-order of magnitude increase in the cost of the trip, so a 500% increase in the cost the trip.

It'd be pretty absurd to try to build the city on Mars if your spaceship just kept staying on Mars not going back to Earth. You'd have this like massive graveyard of ships. You'd have to like do something with them. So it really wouldn't make sense to leave your spaceships on Mars.

You really want to build a propellant plant on Mars and send the ships back. And Mars happens to work out well for that, because it has a CO<sub>2</sub> atmosphere, it's got water ice in the soil, and with H<sub>2</sub>O and CO<sub>2</sub> you can do CH<sub>4</sub> methane and oxygen, O<sub>2</sub>.

	$C_{12}H_{22.4}/O_2$ KEROSENE	$H_2/O_2$ HYDROGEN/OXYGEN	$CH_4/O_2$ DEEP-CRYO METHALOX
VEHICLE SIZE	●	●	●
COST OF PROP	●	●	●
REUSABILITY	●	●	●
MARS PROPELLANT PRODUCTION	×	●	●
PROPELLANT TRANSFER	●	●	●
● GOOD ● OK ● BAD × VERY BAD			

Picking the right propellant is also important.

Think of this as maybe there's three main choices, and they have their merits. But kerosene or rocket-propellant grade kerosene, which is also what jets use — rockets use a very expensive form, a highly refined form of jet fuel, essentially, which is a form of kerosene. It helps keep the vehicle size small, but because it's a very specialized form of jet fuel, it's quite expensive. The reusability potential is lower. Very difficult to make this on Mars, because there's no oil. So



really quite difficult to make propellants on Mars, and then propellant transfer is pretty good but not great.

Hydrogen, although it has a high specific impulse, is very expensive. Incredibly difficult to keep from boiling off, because liquid hydrogen is very close to absolute zero as a liquid. So the insulation required is tremendous, and the energy cost on Mars of producing and storing hydrogen is very high. So we look at the overall system optimization, it was clear to us that methane actually was the clear winner. So it would require maybe anywhere from 50% to 60% of the energy on Mars to refill the propellants, using a propellant depot, and just the technical challenges are a lot easier.

So we think methane is actually better, on, really, almost across the board. And we started off initially thinking that hydrogen would make sense, but we came to the conclusion that the best way to optimize the cost-per-unit mass to Mars and back is to use an all-methane system, technically a deep-cryo methalox.


So those are the four elements that need to be achieved. So whatever architecture, whatever system is designed, whether by SpaceX or anyone, we think these are the four features that need to be addressed in order for the system to really achieve a low cost per trip to the surface of Mars.

And this is a simulation of the overall system. [https://youtu.be/0qo78R\\_yYFA](https://youtu.be/0qo78R_yYFA)

So what you saw there is really quite close to what we will actually build. It will look almost exactly what you saw, so this is not an artist's impression. These... The simulation was actually made from the space engineering CAD (computer-aided design) models.

So this is not, you know, it's not just "this is what it might look like," this is what we plan to try and make it look like.





So in the video you got a sense for what the system architecture looks.


The rocket booster and the spaceship take off, loads the spaceship into orbit. The rocket booster then comes back — it comes back quite quickly, within about 20 minutes — and so it can actually launch the tanker version of the spacecraft, which is essentially the same as the spaceship, but filling up the unpressurized and pressurized cargo areas with propellant tanks.

So they look almost identical, this also helps lower the development costs, which absolutely will not be small. And then the propellant tanker goes up — and it will actually go up multiple times, anywhere from three to five times — to fill the tanks of the spaceship in orbit. And then once the spaceship tanks are full, the cargo has been transferred, and we reach the Mars rendezvous timing, which as I mentioned is roughly every 26 months, that's when the ship would depart.

Now over time there would be many spaceships. Ultimately, I think, upwards of 1,000 or more spaceships waiting in orbit. And so the Mars colonial fleet would depart en masse, kind of "Battlestar Galactica" — if you've seen that thing, it's a good show — so a bit like that. But it actually makes sense to load the spaceships into orbit, because you've got 2 years to do so, and then make frequent use of the booster and the tanker to get really heavy reuse out of those.

And then with the spaceship you get less reuse because you have say, "well, how long is it going to last?" Well, maybe 30 years. So that might be 12, maybe 15 flights of the spaceship, at most. So you really want to maximize the cargo of the spaceship and reuse the booster and the tanker a lot. So the ship goes to Mars, gets propellant replenished, and then returns to Earth.

So I'll go into some of the details of the vehicle design and performance, and I'm going to gloss over — I'll only talk a little bit about the technical details in the actual presentation, and then I'll leave the detailed technical questions to the Q&A that follows.




This is to give us a sense of size. It's quite big.

And the funny thing is I think in the long term, the ships will be even bigger than this. I think that this will represent, this will be relatively small compared to the Mars interplanetary ships of the future.

But it kind of needs to be about this size, because if, in order to fit 100 people thereabouts in the pressurized section, plus carry the luggage and all of the unpressurized cargo — to build propellant plants and build everything from iron foundries, to pizza joints, to you name it — we need to carry a lot of cargo. So it really needs to be roughly on this order magnitude. Because if we say, like, the minimum threshold for self-sustaining city on Mars or civilization would be a million people, well, and you can only go every 2 years — if you have 100 people per ship, that's 10,000 trips.

So I think at least 100 people per trip is the right order of magnitude, and I think we actually may end up expanding the crew section and ultimately taking more like 200 or more people per flight in order to reduce the cost per person. So, 10,000 flights is a lot of flights, so you really want to ultimately think on the order of 1,000 ships.

It will take awhile to build up to 1,000 ships. And so I think if you say, when would we reach that million-person threshold, from the point at which the first ship goes to Mars, it's probably somewhere between 20 to 50 total Mars rendezvous. So it's probably somewhere between maybe 40 to 100 years to achieve a fully self-sustaining civilization on Mars.




So that's the sort of cross-section of the ship, and in some ways it's not that complicated, really.

It's made primarily of an advanced carbon-fiber. The carbon-fiber part is tricky when dealing with deep cryogenics, and trying to achieve both liquid and gas impermeability, and not have gaps occur due to cracking or pressurization that would make the carbon fiber leaky. So this is a fairly significant technical challenge to make deeply cryogenic tanks out of carbon fiber, and it's only recently that we think the carbon fiber technology has gotten to the point where we can actually do this without having to create a liner — some sort of metal liner, or other liner, on the inside of the tanks, which would add mass and complexity.

It's particularly tricky for the hot, gaseous oxygen pressurization. So this is designed to be autogenously pressurized, which means that the fuel and the oxygen, we gassify them through heat exchangers in the engine, and use that to pressurize the tanks. So we'll gassify the methane, and use that to pressurize the fuel tank. Gassify the oxygen, use that to pressurize the oxygen tank.

And this compares — this is a much simpler system than what we have with Falcon 9, where we use helium for pressurization, and we use nitrogen for gas thrusters. In this case we're autogenously pressurized, and then use gaseous methane and oxygen for the control thrusters.

So really you only need two ingredients for this, as opposed to four, in the case of Falcon 9. And actually five, if you consider the ignition liquid we use. It's sort of a complicated liquid to ignite the engines that isn't very reusable. In this case we would use spark-ignition.



So this gives you a sense of vehicles by performance, so current and historic. I don't know if you can actually read that, but: In expendable mode the vehicle that we were proposing, we'd do about 550 tons, and about 300 tons in reusable mode. That compares to Saturn V's max capability of 135 tons.

But I think this really gives a better sense of things.

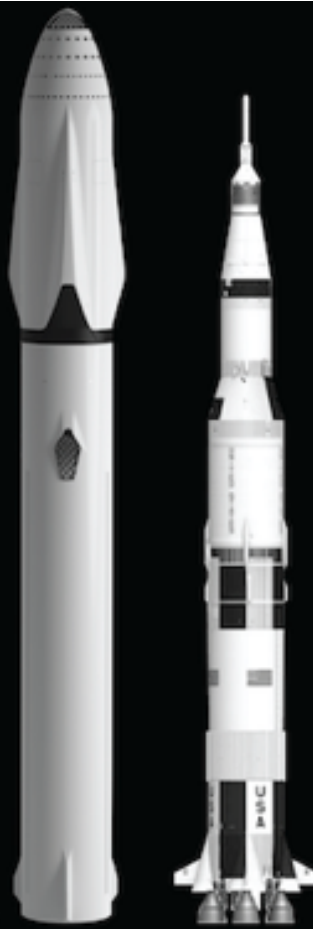
The white bars show the performance of the vehicle. Like, in other words, the payload-to-orbit of the vehicle. So you can see, essentially, what it represents is: What is the size efficiency of the vehicle? And most rockets, including ours — ours that are currently flying — the performance bar is only a small percentage of the actual size of the rocket.

But with the interplanetary system, which will initially be used for Mars, we've been able to — we believe — massively improve the design performance. So it's the first time a rocket's sort of "performance bar" will actually exceed the physical size of the rocket.

This gives you a more direct sort of comparison.

This is, the thrust is quite enormous. We're talking about a liftoff thrust of 13,000 tons. So it will be quite tectonic when it takes off. But it does fit on pad 39A, which NASA has been kind enough to allow us to use, where they somewhat oversized the pad in doing Saturn V. And, as a result, we can actually do a much larger vehicle on that same launch pad. And in the future we expect to add additional launch locations, probably adding one on the south coast of Texas.

	MARS VEHICLE	SATURN V	RATIO
GROSS LIFT-Off MASS(t)	10,500	3,039	3.5
LIFT-Off THRUST (MN)	128	35	3.6
LIFT-Off THRUST (t )	13,033	3,579	3.6
VEHICLE HEIGHT (m)	122	111	1.1
TANK DIAMETER (m)	12	10	1.2
EXPENDABLE LEO PAYLOAD (t)	550	135	4.1
FULLY REUSABLE LEO PAYLOAD (t)	350		




But this gives you a sense of the relative capabilities, if you can read those. But these vehicles have very different purposes. This is really intended to carry huge numbers of people, ultimately millions of tons of cargo to Mars. So you really need something quite large in order to do that.

To talk about some of the key elements of the interplanetary spaceship and rocket booster, we decided to start off the development with what we think are probably the two most difficult elements of the design.

One is the Raptor engine, and this is going to be the highest chamber pressure engine of any kind ever built, and probably the highest thrust-to-weight.

It's a full-flow, stage-combustion engine, which maximizes the theoretical momentum that you can get out of a given source of fuel and oxidizer. We sub-cool the oxygen and methane to densify it. So compared to when propellants are normally used, they're used close to their boiling points in most rockets.



In our case we actually load the propellants close to their freezing point, and that can result in a density improvement of up to around 10% to 12%, which makes an enormous difference in the actual results of the rocket.


It also makes the — it gets rid of any cavitation risk for the turbo pumps, and it makes it easier to feed a high-pressure turbo pump if you have very cold propellant.

Really one of the keys here, though, is the vacuum version of Raptor having a 382-second ISP [specific impulse]. This is really quite critical to the whole Mars mission, and we're confident we can get to that number, or at least within a few seconds of that number, ultimately maybe exceeding it slightly.

The rocket booster, in many ways, is really a scaled-up version of the Falcon 9 booster.

You'll see a lot of similarities, such as the grid fins, obviously clustering a lot of engines at the base, and the big difference really being that the primary structure is an advanced form a

carbon-fiber, as opposed to aluminum-lithium. And that we use autogenous pressurization, and we get rid of the helium and the nitrogen.



So, this uses 42 Raptor engines. It's a lot of engines, but we use nine on a Falcon 9, and with Falcon Heavy, which should launch early next year, there's 27 engines on the base. So we've got pretty good experience with having a large number of engines. It also gives us redundancy, so that if some of the engines fail, you can still continue the mission and be fine.

But the main job of the booster is to accelerate the spaceship to around 8,500 kilometers an hour. For those that aren't as familiar with orbital dynamics, really it's all about velocity and not about height. So really that's the job of the boosters. The booster's like the javelin thrower — so it's gotta toss that javelin, which is the spaceship. In the case of other planets, though, which have a gravity well that is not as deep. So Mars, the moons of Jupiter, see if we went to maybe even Venus — Venus will be a little trickier — but for most of the solar system, you only need the spaceship.

You don't need the booster if you have a lower gravity wells. So no booster is needed on the moon or Mars or any other moons of Jupiter or Pluto. You just need the spaceship. The booster is just there for heavy gravity wells.

And then we've also been able to optimize the propellant needed for boost-back and landing, to get it down to about 7%, of the liftoff propellant load, and we think with some optimization we can get down to about 6%.

And we're also getting quite comfortable with the accuracy of the landing. If you've been watching the Falcon 9 landings, you'll see that they're getting increasingly closer to the bull's-




eye. And we think, particularly with the addition of some thrusters, maneuvering thrusters, we can actually put the booster right back on the launch stand. And then those fins at the base are essentially centering features to take out any minor position mismatch at the launch site.

So this is what it looks like at the base so we think we only need to gimbal or steer the center cluster of engines.

So there's seven engines in the center cluster, those would be the ones that that move, for steering the rocket, and the other ones would be fixed in position, which gives us the best concentration of — we can max out the number of engines because we don't have to leave any room for gimbaling or moving the engine.

And this is all designed so that you could actually lose multiple engines, even at liftoff or anywhere in flight, and continue the mission safely.



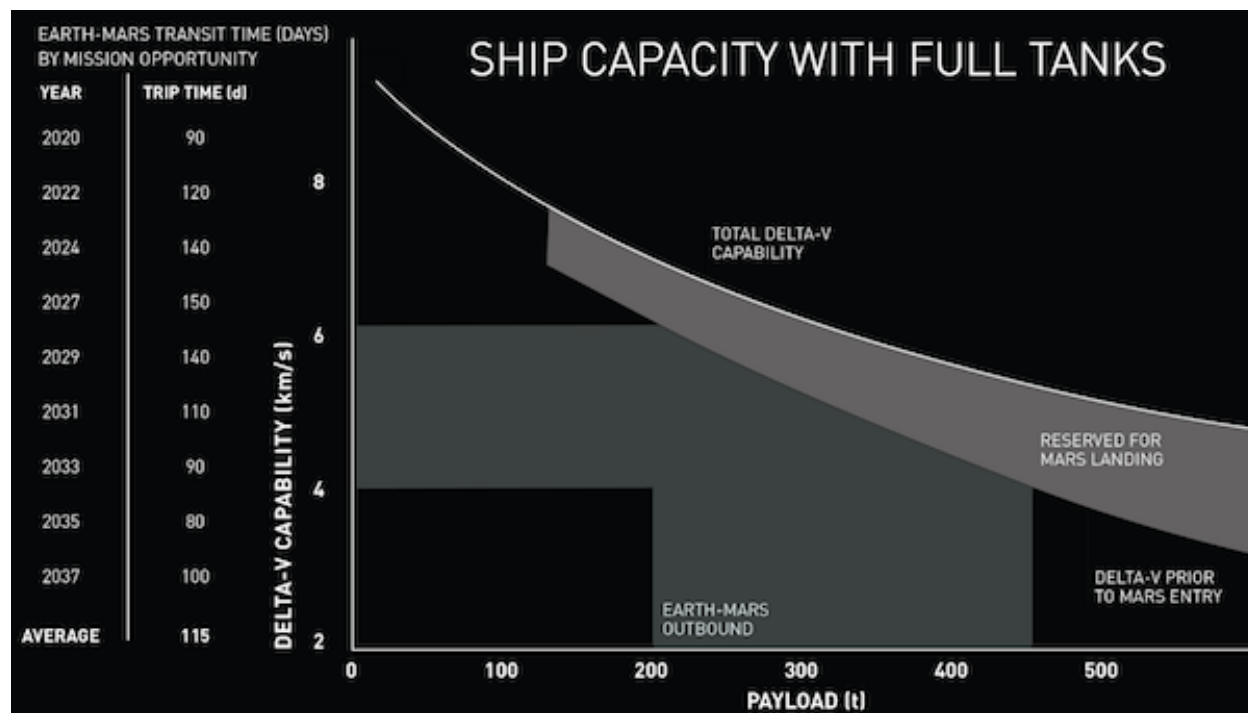
So for the spaceship itself, in the top we have the pressurized compartment — and I'll show you a fly-through of that in a moment — and beneath that is where we would have the unpressurized cargo, which will be really flat-packed, in a very dense format. And below that is the liquid oxygen tank.

The liquid oxygen tank is probably the hardest piece of this whole vehicle because it's gotta handle propellant at the coldest level, and the tanks themselves actually form the airframe. So that the air frame structure and the tank structure are combined, as it is in all modern rockets, and in an aircraft. For example, the wing is really a fuel tank in wing shape. So it has to take the thrust loads of ascent, the loads of re-entry, and then it has to be impermeable to gaseous oxygen, which is tricky, and non-reactive to gaseous oxygen.

So that's the hardest piece of the spaceship itself, which is actually why we started on that element. I will show you some pictures of that later.

And then below the oxygen tank is the fuel tank, and then the engines are mounted directly to the thrust cone on the base. And then there are six of the vacuum, the high-efficiency vacuum engines around the perimeter, and those don't gimbal. And then there are three of the sea-level versions of the engine, which do gimbal and provide the steering. Although we can do some amount of steering, if you're in space, with differential thrust on the outside engines.

The net effect is a cargo-to-Mars of up to 450 tons, depending upon how many refills you do with the tanker. The goal is at least 100 passengers per ship, although I think ultimately we'll probably see that number go to 200 or more.




This chart's a little difficult to interpret at first, but I kind of decided to put it there for people who wanted to watch the video afterwards and sort of take a closer look, analyze some of the numbers.

The column on the left is probably what's most relevant, and that gives you the trip time.


So depending upon which Earth-Mars rendezvous you're aiming for, the trip time at six kilometers per second, departure velocity, can be as low as 80 days. And then, over time, I think we'd obviously improve that, and ultimately I suspect that you'd see Mars transit times of as little as 30 days in the more distant future.

So it's fairly manageable, considering the trips that people used to in the old days. They'd routinely take sailing voyages that would be 6 months or more.

So on arrival the heat shield technology is extremely important.



We've been refining the heat-shield technology using our Dragon spacecraft, and we're now on version three of PICA, which is "phenolic impregnated carbon ablator," and it's getting more and more robust with each new version, with less ablation, more resistance, less need for refurbishment.



The heat shield's basically a giant brake pad. So it's like, how good can you make that brake pad against extreme reentry conditions, and minimize the cost of refurbishment. And make it so that you could have many flights with no refurbishment at all.


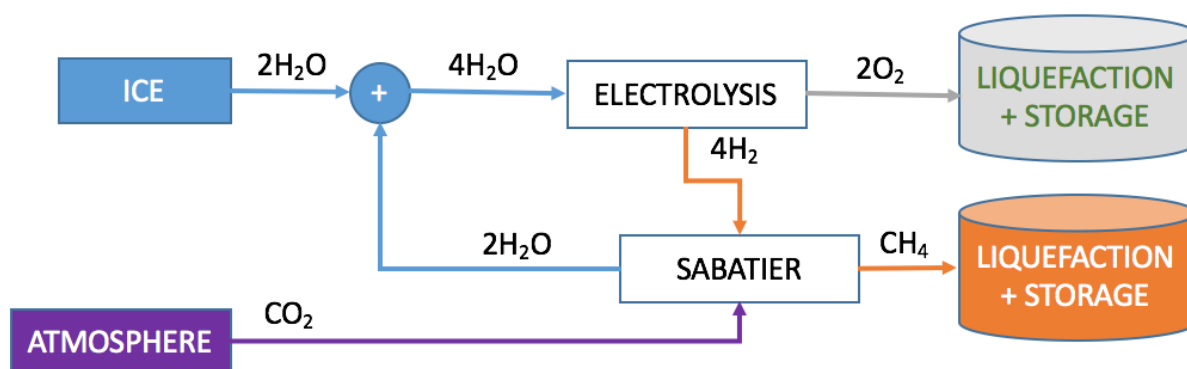
This is a fly-through of the crew compartment. Just want to give you a sense of what it would feel like to actually be in the spaceship. (<https://www.youtube.com/watch?v=PC0c5tFmltQ>)

I mean, in order to make it appealing, and increase that portion of the Venn diagram of people actually wanting to go, it's gotta be really fun and exciting and it can't feel cramped or boring.

So the crew compartment or the occupant compartment, is set up so that you can do zero-g games, you can float around, there'll be like movies, lecture halls, you know, cabins, a restaurant — it will be, like, really fun to go. You're gonna have a great time.

So that propellant plant on Mars.

Again, this is one of the slides that I won't go into detail here, but people can think about offline.



The key point being that the ingredients are there on Mars to create a propellant plant with relative ease, because the atmosphere is primarily  $\text{CO}_2$ , and there's water ice almost everywhere. You've got the  $\text{CO}_2$  plus  $\text{H}_2\text{O}$  to make methane,  $\text{CH}_4$ , and oxygen  $\text{O}_2$ , using the Sabatier reaction.

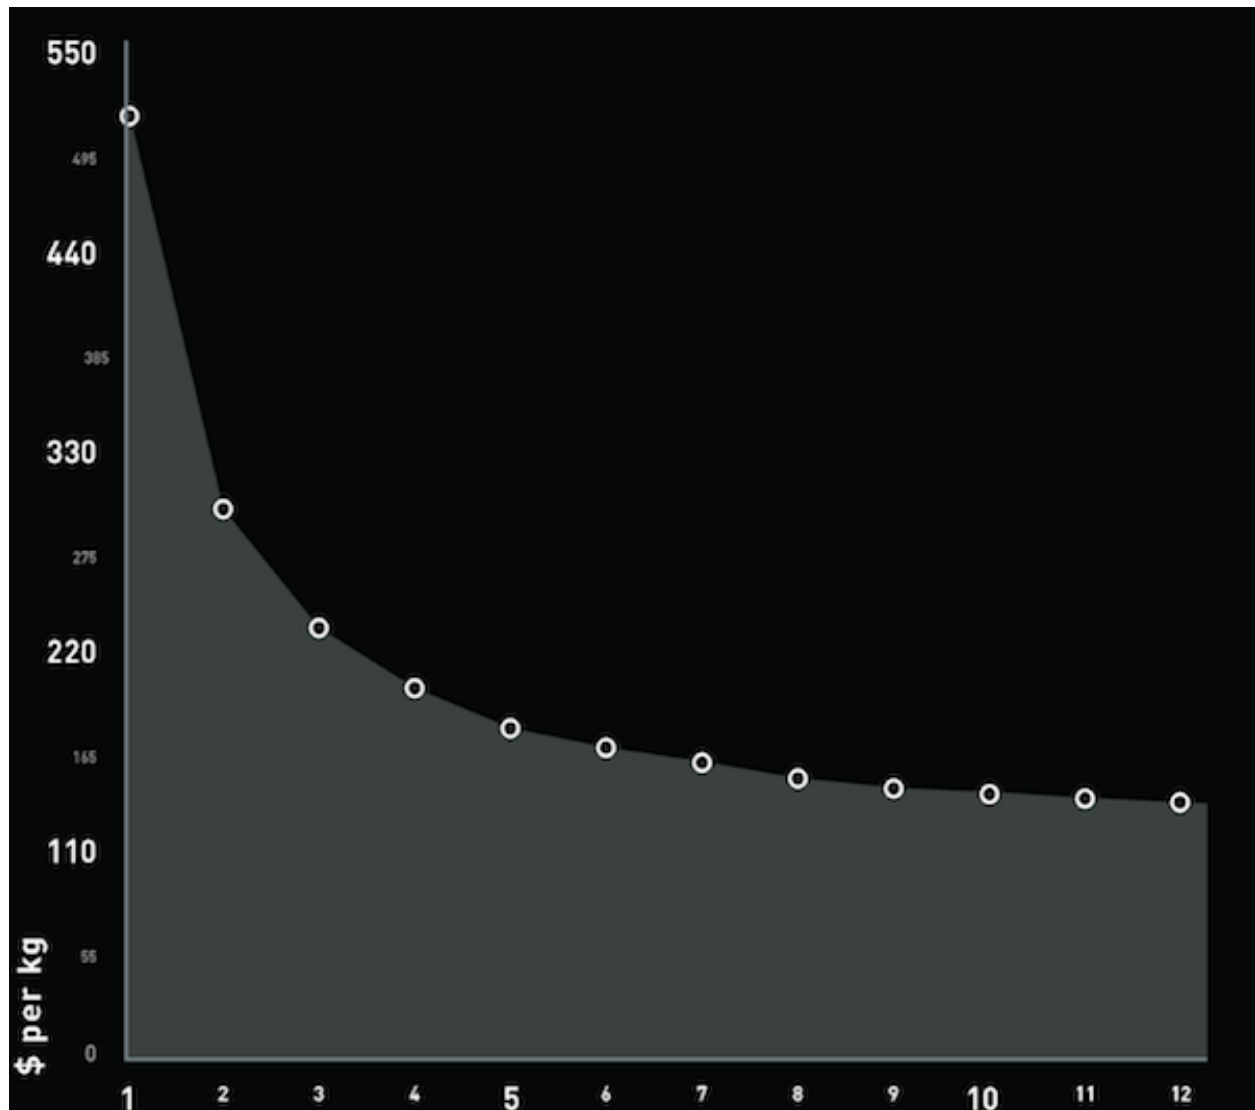
There are effectively unlimited supplies of carbon dioxide and water on Mars; 5 million cubic km ice, 25 trillion metric tons  $\text{CO}_2$ .

The trickiest thing, really, is the energy source, which we think we can do with a large field of solar panels.

	BOOSTER	TANKER	SHIP
FABRICATION COST	\$230M	\$130M	\$200M
LIFETIME LANCHES	1,000	100	12
AMORTIZED FABRICATION COST/LAUNCH	\$0.23M	\$1.3M	\$16.7M
LANCHES PER MARS TRIP	6	5	1
AVERAGE MAINTENANCE COST PER USE	\$0.2M	\$0.5M	\$10M
LAUNCH COST	\$0.2M		
COST OF PROPELLANT (\$168/t) PER LAUNCH	\$0.7M		
ESTIMATED TOTAL COST PER MARS TRIP	\$11M	\$8M	\$43M
CARGO DELIVERED TO MARS PER TRIP FOR \$62M	450T		
COST/TON DELIVERED TO MARS	\$140K		

So then to give you a sense of the cost, really the key is making this affordable to almost anyone who wants to go. And we think, based on this architecture, assuming optimization over time,

like, the very first flights would be fairly expensive — but the architecture allows for a cost-per-ticket of less than \$200,000, maybe as little as \$100,000 over time, depending upon how much mass a person takes.



So we're, right now, estimating about \$140,000 per ton to the surface of Mars. So if a person plus their luggage is less than that, taking into account food consumption and life-support, then we think that the cost of moving to Mars could drop below \$100,000.

So, funding. We've thought about funding sources.

And so it's steal underpants, launch satellites, send cargo to space station, Kickstarter — of course — followed by profit. So obviously it's going to be a challenge to fund this whole endeavor.

We do expect to generate pretty decent net cash flow from launching lots of satellites and servicing the space station for NASA, transferring cargo to and from the space station, and then I know that there's a lot of people in the private sector who are interested in helping fund a base on Mars. And then perhaps there'll be interest from the government sector side to also do that.

Ultimately this is going to be a huge public-private partnership, and I think that's how the United States is established, and in many other countries around the world — is a public-private partnership. So I think that's probably what occurs, and right now we're just trying to make as much progress as we can with the resources that we have available, and just sort of keep moving both forward, and hopefully I think, as we show that this is possible, that this dream is real, not just a dream — it's something that can be made real — I think the support will snowball over time.

And I should say also that the main reason I'm personally accumulating assets is in order to fund this. So I really don't have any other motivation for personally accumulating assets, except to be able to make the biggest contribution I can to making life multi-planetary.

Timelines. I'm not the best at this sort of thing.

But just to show you where we started off, in 2002 SpaceX basically consisted of carpet and a mariachi band. That was it. That's all of SpaceX in 2002. As you can see I'm a dancing machine, and yeah I believe in kicking off celebratory events with mariachi bands. I really like mariachi bands.






But that was what we started off with in 2002, and really, I thought maybe we had a 10% percent chance of doing anything — of even getting a rocket to orbit, let alone getting beyond that and taking Mars seriously.

But I came to the conclusion that if there wasn't some new entrant into the space arena with a strong ideological motivation, then it didn't seem like we were on a trajectory to ever be a space-faring civilization and be out there among the stars. Because in '69 we were able to go to the moon, and the space shuttle could get to low-earth orbit, and then obviously the space shuttle got retired, but that trend line is down to zero.

So I think what a lot of people don't appreciate is that technology does not automatically improve. It only improves if a lot of really strong engineering talent is applied to the problem that it improves. And there are many examples in history where civilizations have reached a certain technology level, and then have fallen well below that and then recovered only millennia later.



So we go from 2002, where we're basically clueless, and then with Falcon 1 — the smallest useful orbital rocket that we could think of, which would deliver half a ton to orbit. And then 4 years later we developed, we built the first vehicle.



So we developed the main engine, the upper-stage engine, the airframes, the fairing, and the launch system, and had our first attempt at launch in 2006, which failed. So, that lasted about 60 seconds, unfortunately.

But it was 2006, 4 years after starting, which is also when we actually got our first NASA contract. And I just want to say that I'm incredibly grateful to NASA for supporting SpaceX, despite the fact that our rocket crashed. It was awesome, I'm NASA's biggest fan — so “thank you very much” to the people that had the faith to do that.

So then 2006, followed by a lot of grief, and then finally the fourth launch of Falcon 1 worked in 2008. And we were really down to our last pennies. In fact, I only thought I had enough money for three launches, and the first three bloody failed, and we were able to scrape together enough to just barely make it into a fourth launch, and that — thank goodness — that fourth launch succeeded in 2008.

That was a lot of pain.

And then also at the end of 2008 is when NASA awarded us the first major operational contract, which was for resupplying cargo to the space station and bringing cargo back. Then a couple of years later we did the first launch of Falcon 9, version 1, and that had about a 10-ton-to-orbit capability, so it was about 20 times the capability of Falcon 1, and also assigned to carry our Dragon spacecraft.

Then 2010 is our first mission to the space station, so we were able to finish development of Dragon and dock with the space station in 2010. Sorry — 2010 is expendable Dragon, 2012 is when we delivered and returned cargo from the space station.

2013 is when we first started doing vertical takeoff and landing tests.

And 2014 is when we were able to have the first orbital booster do a soft landing in the ocean. The landing was soft, then it fell over and exploded, but the landing — for 7 seconds — it was good. And we also improved the capability of the vehicle from 10 tons to about 13 tons to LEO (Low Earth orbit).

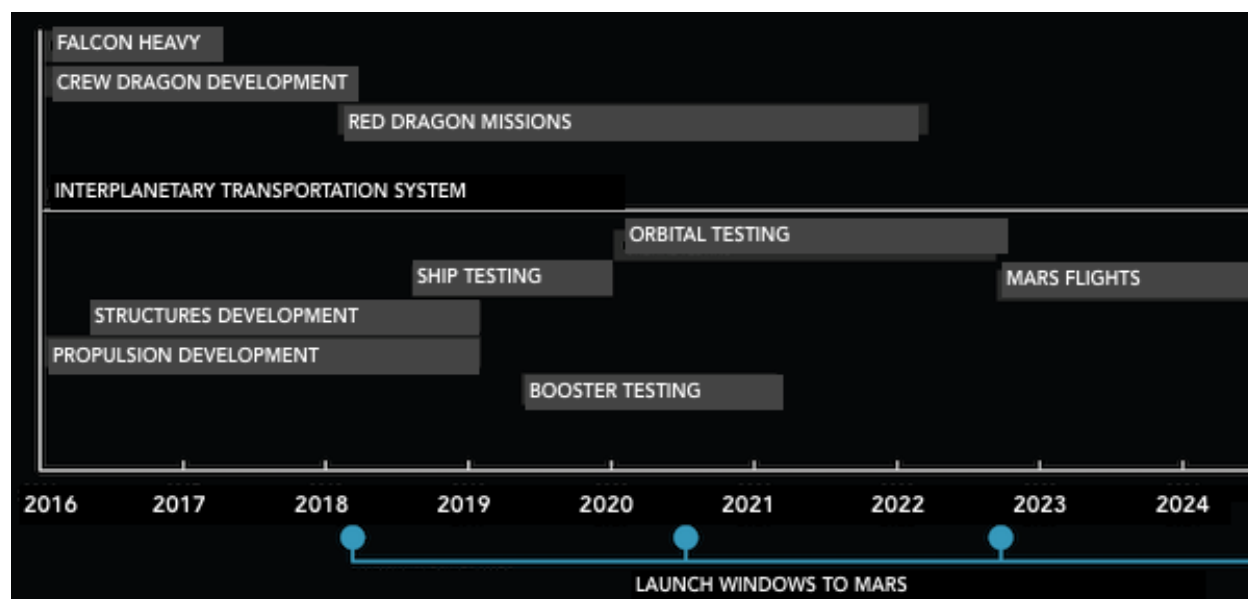
And then 2015, last year in December, that was definitely one of the best moments of my life: when the rocket booster came back and landed at Cape Canaveral. That was really... yeah.

So I think that really showed we could bring an orbit-class booster back from a very high velocity, all the way to the launch site, and land it safely, and with almost no refurbishment required for re-flight. And if things go well, we are hoping to re-fly one of the landed boosters in a few months.

So, yeah, and then 2016 we also demonstrated landing on a ship. The landing on the ship is very important for very high-velocity geosynchronous missions, and that's important for reusability

of Falcon 9 because roughly a quarter of our missions are sort of servicing the space station, and then there's a few other low-Earth orbit missions. But most of our missions — probably 60% of our missions — are commercial geo [geosynchronous] missions. So we've got to do these high-velocity missions that really need to land on the ship out to sea. They don't have enough propellant on board to boost back to the launch site.

So looking at the future, next steps.



We were kind of intentionally a bit fuzzy about this timeline. But the... We're going to try to make as much progress as we can, obviously it's a very constrained budget, but we're going to try to make as much progress as we can on the elements of the interplanetary transport booster and spaceship. And hopefully we'll be able to do, to complete the first development spaceship in maybe about 4 years, and start doing suborbital flights with that.

In fact, it actually has enough capability that you could maybe even go to orbit if you limit the amount of cargo with the spaceship. But you have to really, just have to really strip it down. But in tanker form it can definitely get to orbit. It can't get back, but we can get to orbit.

Actually, I was sort of thinking, like, maybe there is some sort of market for really fast transport of stuff around the world, provided we can land somewhere where noise is not a super-big deal — rockets are very noisy — but we could transport cargo to anywhere on earth in 45 minutes, at the longest. So most places on Earth would be maybe 20, 25 minutes. So maybe if we had a floating platform out off the coast of the USA, off the coast of New York, say 20 or 30 miles out, you could go from, you know, from New York to Tokyo in — I don't know — 25 minutes. Cross the Atlantic in 10 minutes. Really, most of your time would be getting to the ship. And then it'd be real quick after that.

So there's some intriguing possibilities there, although we're not counting on that.

And then development of the booster. And actually the booster part is relatively straightforward, because it amounts to a scaling-up of the Falcon 9 booster. So there's, we don't see a lot of showstoppers there. So yeah.

But then trying to put it all together and make this actually work for Mars. If things go super-well, it might be kind of in the 10-year time frame. But I don't wanna say that's when it will occur, there's a huge amount of risk, it's going to cost a lot, good chance we won't succeed, but we're going to do our best, and we're going try to make as much progress as possible.

Oh and we're gonna try to send something to Mars on every Mars rendezvous from here on out. So Dragon 2, which is a propulsive lander, we plan to send to Mars in a couple years. And then do probably another Dragon mission in 2020.



In fact, we want to establish a steady cadence — that there's always a flight leaving, like a train leaving the station. With every Mars rendezvous we will be sending a Dragon — at least a Dragon to Mars, and ultimately the big spaceship — so if there are people that are interested in putting payloads on Dragon, you know you can count on a ship that's going to transport something on the order of at least 2 or 3 tons of useful payload to the surface of Mars.

Yes, that's part of the reason why we designed Dragon 2 to be a propulsive lander.

As a propulsive lander, you can go anywhere in the solar system. So you could go to the moon, you could go to... Well, anywhere, really. Whereas if something relies on parachutes or wings, then you can pretty much only — well if it's wings, you can pretty much only land on Earth, because you need a runway, and most places don't have a runway. And then any place that doesn't have a dense atmosphere, you can't use parachutes.

But propulsive works anywhere. So Dragon should be capable of landing on any solid or liquid surface in the inner solar system.



And then I was real excited to see that the team managed to do the, all of our Raptor engine firing, in advance of this conference. I just want to say thanks to the Raptor team for really working 7 days a week to try to get this done in advance of the presentation, because I really want to show that we've made some hardware progress in this direction, and the Raptors are really tricky engines.

It's a lot trickier than the Merlin, because it's a full-flow stage combustion — much higher pressure. I'm kind of amazed it didn't blow up on the first firing, but fortunately it was good.

Yeah. It's kind of interesting to see the “mock diamonds” forming.

So the... Part of the reason for making engines small... Raptor, although it has three times the thrust of a Merlin, is actually only about the same size as the Merlin engine, because it has three times the operating pressure. And that means we can use a lot of the production techniques that we've honed with Merlin.

We're currently producing Merlin engines at almost 300 per year. So we understand how to make rocket engines in volume, so even though the launch vehicle uses 42 on the base and nine on the upper stage — so we have 51 engines to make — that's well within our production capabilities for Merlin, and this is a similarly sized engine to Merlin, except for the expansion ratio, so we feel really comfortable about being able to make this engine in volume at a price that doesn't break our budget.

And then we also wanted to make progress on the primary structure, so as I mentioned this is really... a very difficult thing to make. Is to make something out of carbon fiber.

Even though carbon fiber has incredible strength-to-weight, when you want one of them put super-cold liquid oxygen and liquid methane — particularly liquid oxygen — in the tank, it's subject to cracking and leaking and it's a very difficult thing to make.

Just the sheer scale of it is also challenging, because you've gotta lay out the carbon fiber in exactly the right way on a huge mold, and you've gotta cure that mold at temperature, and then it's... just really hard to make large carbon-fiber structures that can do all of those things and carry incredible loads.

So that's the other thing we wanted to focus on, was the Raptor, and then building the first development tank for the Mars spaceship.



So this is really the hardest part of the spaceship. The other pieces are, we have a pretty good handle on, but this was the trickiest one. So we wanted to tackle it first. You get a sense for how big the tank is. It's really quite big.

Also, big congratulations to the team that worked on that — they were also working seven days a week to try to get this done in advance of the IAC.



And so we managed to build the first tank and initial tests with the cryogenic propellant actually look quite positive. We have not seen any leaks or major issues.

This is what the tank looks like on the inside.




So you can get a real sense for how much, just how big this tank is. It's actually completely smooth on the inside, but the way that the carbon-fiber plys reflect the light makes it look faceted.

So what about beyond Mars?

So as we thought about this system, and the reason we call it a system — because generally I don't like calling things systems, because everything's a system, including your dog — is that it's actually more than a vehicle. There's this rocket booster, the spaceship, the tanker, and the propellant plant, the *in situ* propellant production.

If you have all of those four elements, you can actually go anywhere in the solar system by planet-hopping or by moon-hopping.

So by establishing a propellant depot in the Asteroid Belt, or on one of the moons of Jupiter, you can go, you can make flights from Mars to Jupiter — no problem.



In fact, even from, even without a propellant depot at Mars, you could do a flyby of Jupiter without a propellant depot.

But by establishing a propellant depot, let's say, you know, Enceladus or Europa, or any — there's a few options — and then doing another one on Titan, Saturn's moon, and then perhaps another one further out on Pluto, or elsewhere in the solar system...






This system really gives you freedom to go anywhere you want in the greater solar system. So you could actually travel out to the Kuiper Belt, the Oort Cloud.

I wouldn't recommend this for interstellar journeys, but this basic system — provided we have filling stations along the way — is, means full access to the entire greater solar system.

It'd be really great to do a mission to Europa, particularly.



Alright so, any questions that I can answer?